CHAPTER 5

PHYSICAL AND LOGICAL TOPOLOGIES

After reading this chapter and completing the exercises, you will be able to:

- Describe the basic and hybrid LAN physical topologies, their uses, advantages, and disadvantages
- Describe a variety of enterprise-wide and WAN physical topologies, their uses, advantages, and disadvantages
- Compare the different types of switching used in data transmission
- Understand the transmission methods, or logical topologies, underlying Ethernet, Token Ring, LocalTalk, and FDDI networks

ON THE JOB

I've worked as a senior networking engineer at a large furniture manufacturer for almost a decade. Many years ago, a consultant had nearly convinced us that we should upgrade our old Thinnet LAN to Token Ring. He insisted that Token Ring was much more reliable, scalable, and easier to support. This was true enough, compared to our old network. But replacing the cabling, connectivity devices, and NICs (not to mention retraining our staff to use the new network) would be a huge investment, so our IT department took plenty of time to mull it over. The other senior engineer and I had an ongoing debate about whether Token Ring or Ethernet would win out as the preferred standard of the networking world. He pointed out that Token Ring, although slightly more expensive to implement, was more efficient. Therefore, he projected, it could surpass Ethernet in speed, given the right advances in technology. I had a feeling that Ethernet, because it was less expensive and gaining in popularity, would win out as the transmission of choice.

In the end, we all agreed to replace our Thicknet network with Token Ring. It took almost a full year to complete the transition, and we all quickly learned about the idiosyncrasies of this unfamiliar transmission method. Unfortunately, however, a few years later it became clear that Ethernet technology was quickly advancing and becoming the standard of the networking world. Fast Ethernet was going to be a reality. Meanwhile, our Token Ring equipment was becoming more expensive and seemed too slow for our needs. We recognized that although Token Ring is a solid technology, it might be headed the way of the dinosaurs. We began the process of evaluating our network all over again, and in 1999 decided to make the plunge to Fast Ethernet.

Stan Myzlowski
Hartford Products

Just as an architect of a house must decide where to place walls and doors, where to install electrical and plumbing systems, and how to manage traffic patterns through rooms to make a house more livable, a network architect must consider many factors, both seen and unseen, when designing a network. This chapter details some basic elements of network architecture: physical and logical topologies. These elements are crucial to understanding networking hardware, design, troubleshooting, and management, all of which are discussed later in this book.

In this chapter you will learn the basic layouts of LANs and WANs, the advantages and disadvantages of each layout, and the optimal application for each layout. You will also learn about the most commonly used logical topologies, or network transmission methods: Ethernet, Token Ring, LocalTalk, FDDI, and ATM. Once you master the physical and logical fundamentals of network architecture, you will have all the tools necessary to design a network as elegant as the Taj Mahal.

Simple Physical Topologies

A **physical topology** is the physical layout, or pattern, of the nodes on a network. It depicts a network in broad scope; that is, it does not specify device types, connectivity methods, or addresses on the network. Physical topologies are divided in three fundamental geometric shapes: bus, ring, and star. These shapes can be mixed to create hybrid topologies. Before you design a network, you need to understand physical topologies, because they can affect which logical topology you use (for example, Ethernet or Token Ring), how your building is cabled, and what kind of network media you use. You must also understand a network's physical topology to troubleshoot its problems or change its infrastructure. A thorough knowledge of physical topologies is necessary to obtain Network+ certification.

This chapter builds on the terms and concepts discussed in Chapters 2, 3, and 4. If you do not have a clear understanding of the material covered in those chapters, take time to review them now.

Physical topologies and logical topologies (discussed later) are two different networking concepts. You should be aware that when used alone, the word "topology" often refers to a network's *physical* topology.

Bus

A **bus topology** consists of a single cable connecting all nodes on a network without intervening connectivity devices. Figure 5-1 depicts a typical bus topology.

The single cable is called the **bus** and can support only one channel for communication; as a result, every node shares the bus's total capacity. Most bus networks—for example, Thinnet and Thicknet—use coaxial cable as their physical medium. A bus topology can be considered a peer-to-peer topology, because every device on the network shares the responsibility for getting data from one point to another. Each node on a bus network passively listens for data directed to it. When one node wants to transmit data to another node, it broadcasts an alert to the entire network, informing all nodes that a transmission is being sent; the destination node then picks up the transmission. Nodes between the sending and receiving nodes ignore the message.

For example, suppose that you want to send an instant message to your friend Diane, who works across the hall, asking whether she wants to have lunch with you. You click the send button after typing your message, and the data stream that contains your message is sent to your NIC. Your NIC then sends a message across the shared wire that essentially says, "I have a message for Diane's computer." The message passes by every NIC between your computer and Diane's computer until Diane's computer recognizes that the message is meant for it and responds by accepting the data.

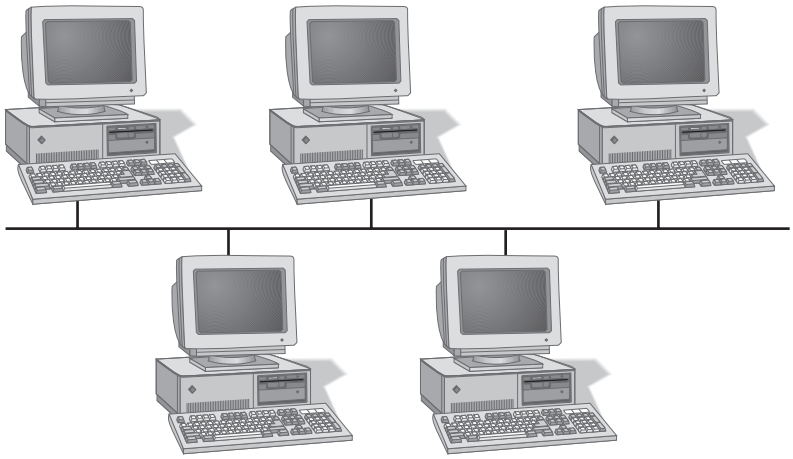

Figure 5-1 A bus topology network

At the ends of each bus network are 50-ohm resistors known as terminators. As you learned in Chapter 4, terminators stop signals after they have reached the end of the wire. Without these devices, signals on a bus network would travel endlessly between the two ends of the network—a phenomenon known as signal bounce—and new signals could not get through. To understand this concept, imagine that you and a partner, standing at opposite sides of a canyon, are yelling to each other. When you call out, your words echo; when your partner replies, his words also echo. Now imagine that the echoes never fade. After a short while you could not continue conversing because all of the previously generated sound waves would still be bouncing around, creating too much noise for you to hear anything else. On a network, terminators prevent this problem by halting the transmission of old signals. Figure 5-2 depicts a terminated bus network.

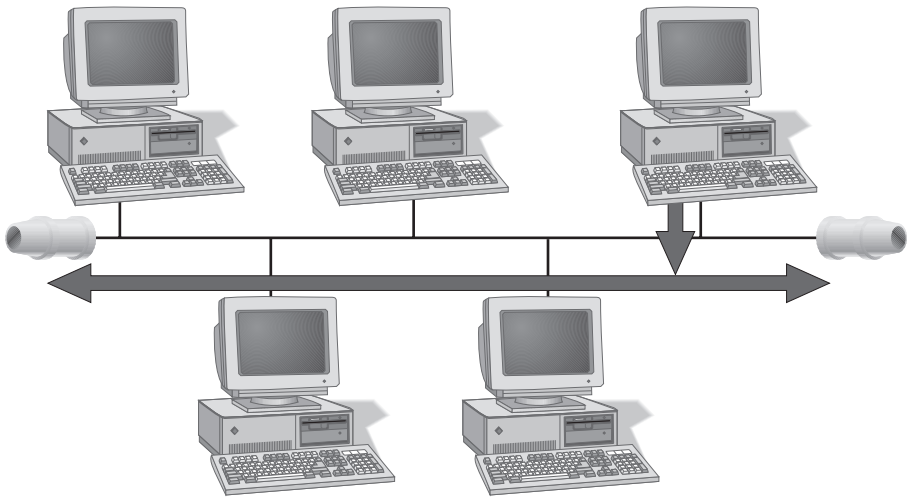

Figure 5-2 A terminated bus network

Although networks based on a bus topology are inexpensive to set up, they do not scale well. As you add more nodes, the network's performance degrades. Because of the single-channel limitation, the more nodes on a bus network, the more slowly the network will transmit and deliver data. For example, suppose a bus network in your small office supports two workstations and a server, and saving a file to the server takes two seconds. During that time, your NIC first checks the communication channel to make sure it is free, then issues data directed to the server. When the data reach the server, the server accepts them. Suppose, however, that your business experiences tremendous growth, and you add five more workstations during one weekend. The following Monday, when you attempt to save a file to the server, the save process might take five seconds, because the new workstations may also be using the communications channel, and your workstation may have to wait for a chance to transmit. As this example illustrates, a bus topology would not be practical for a network of more than 200 workstations. In fact, it is rarely practical for networks with more than a dozen workstations.

Bus networks are also difficult to troubleshoot, because it is a challenge to identify fault locations. To understand why, think of the game called "telephone," in which one person whispers a phrase into the ear of the next person, who whispers the phrase into the ear of another person, and so on, until the final person in line repeats the phrase aloud. The vast majority of the time, the phrase recited by the last person bears little resemblance to the original phrase. When the game ends, it's hard to determine precisely where in the chain the individual errors cropped up. Similarly, errors may occur at any intermediate point on a bus network, but at the receiving end it's possible to tell only that an error occurred. Finding the source of the error can prove very difficult, because you cannot retrace the data's progress from one node to the next; that is, the nodes don't "remember" the data after they pass it on. (In the telephone game analogy, this situation would be similar to every person in the line forgetting the phrase after he or she passed

it on—a situation that would make it impossible to trace the evolution of the phrase as it moves from one person to the next.)

A final disadvantage to bus networks is that they are not very fault-tolerant, because a break or a defect in the bus affects the entire network. As a result, and because of the other disadvantages associated with this topology, you will rarely see a network run on a pure bus topology. You may, however, encounter hybrid topologies that include a bus component, as discussed later in this chapter.

Ring

In a **ring topology**, each node is connected to the two nearest nodes so that the entire network forms a circle, as shown in Figure 5-3. Data are transmitted clockwise, in one direction (unidirectionally), around the ring. Each workstation accepts and responds to packets addressed to it, then forwards the other packets to the next workstation in the ring. Because a ring network has no "ends," and because data stop at their destination, ring networks do not require terminators. In most ring networks, twisted-pair or fiber-optic cabling is used as the physical medium.

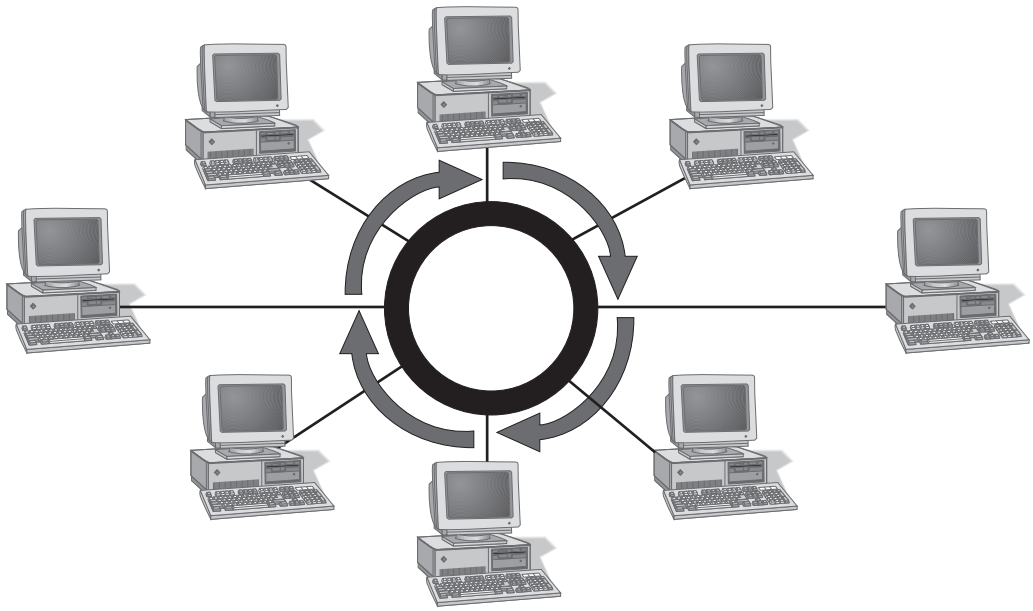

Figure 5-3 A typical ring network

One method for passing data on a ring network is token passing. In **token passing**, a 3-byte packet, called a token, is transmitted from one node to another around the ring. If a computer on the ring has information to transmit, it picks up the token packet, adds control and data information plus the destination node's address to transform the token

into a data frame, and then passes the token on to the next node. The transformed token, now in the form of a frame, circulates around the network until it reaches its intended destination. The destination node picks it up and returns an acknowledgment message to the originating node. After the originating node receives the acknowledgment, it releases a new free token and sends it down the ring. This approach ensures that only one workstation transmits data at any given time. Because each workstation participates in sending the token around the ring, this architecture is known as an **active topology**. Each workstation acts as a repeater for the transmission.

The drawback of a simple ring topology is that a single malfunctioning workstation can disable the network. For example, suppose that you and five colleagues share a pure ring topology LAN in your small office. You decide to send an instant message to Thad, who works three offices away, telling him that you accidentally received a package addressed to him. Between your office and Thad's office are two other offices, and two other workstations on the ring. Your instant message must pass through the two intervening workstations' NICs before it reaches Thad's computer. If one of these workstations has a malfunctioning NIC, your message will never reach Thad.

In addition, just as in a bus topology, the more workstations that must participate in token passing, the slower the response time. Consequently, pure ring topologies are not very flexible or scalable.

Contemporary LANs rarely use pure ring topologies. A variation of the ring topology, known as a star-wired ring, is popular for some types of networks, such as Token Ring networks. Star-wired rings and Token Ring technology will be discussed later in this chapter.

Star

In a **star topology**, every node on the network is connected through a central device, such as a hub. Figure 5-4 depicts a typical star topology. Star topologies are usually built with twisted-pair or fiber cabling. Any single cable on a star network connects only two devices (for example, a workstation and a hub), so a cabling problem will affect two nodes at most. Devices such as workstations or printers transmit data to the hub, which then retransmits the signal to the network segment containing the destination node.

Star topologies require more cabling than ring or bus networks. They also require more configuration. However, because each node is separately connected to a central connectivity device, they are more fault-tolerant. A single malfunctioning cable or workstation cannot disable star networks. A failure in the central connectivity device (such as a hub), can take down a LAN segment, though.

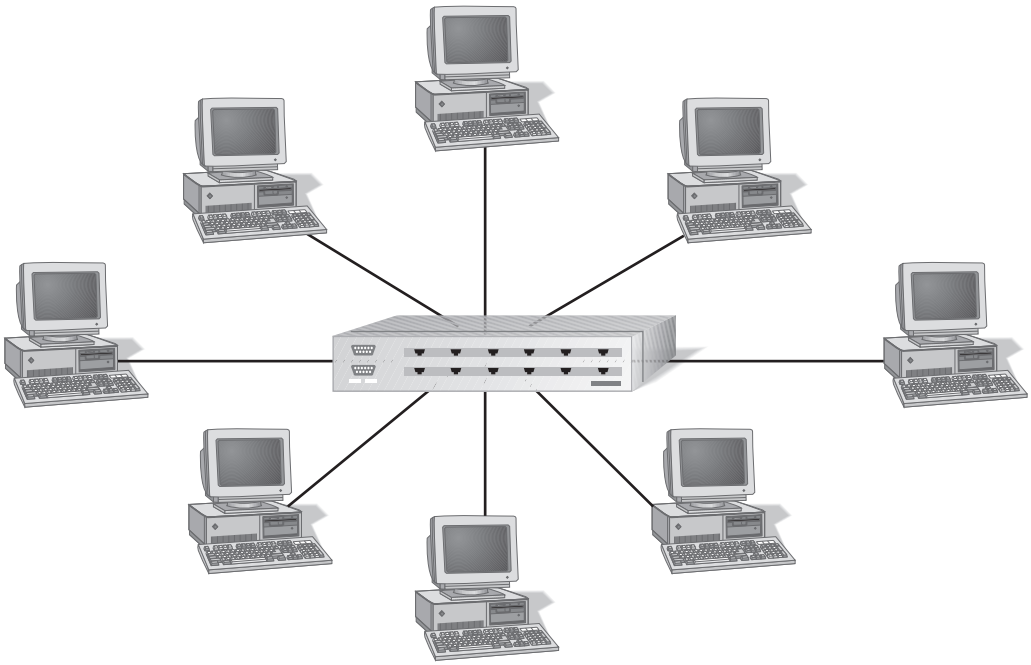

Figure 5-4 A typical star topology network

Because they include a centralized connection point, star topologies can easily be moved, isolated, or interconnected with other networks; they are therefore scalable. For this reason, and because of their fault tolerance, the star topology has become the most popular fundamental layout used in contemporary LANs. Many network administrators have replaced their old bus or ring networks with star networks in recent years. Single star networks are commonly interconnected with other networks through hubs and switches to form more complex topologies.

As you learned in Chapter 4, 10BaseT and 100BaseT Ethernet networks are based on the star topology, as are most LocalTalk networks. These networks can support a maximum of only 1024 addressable nodes on a logical network. Thus, you can say that star networks support a maximum of 1024 nodes. For example, if you have a campus with 3000 users, hundreds of networked printers, and scores of other devices, you must strategically create smaller logical networks. Even if you had 1000 users and *could* put them on the same logical network, you probably wouldn't, because doing so would result in poor performance and impossible management. Instead, you should subdivide the users and their peripherals into workgroups according to their needs or geographic locations. Chapter 16 describes the process of evaluating user and organizational requirements when designing a network.

Hybrid Physical Topologies

Except in very small networks, you will rarely encounter a network that follows a pure bus, ring, or star topology. Simple topologies are too restrictive, particularly if the LAN must accommodate a large number of devices. More likely, you will work with a complex combination of these topologies, known as a **hybrid topology**. Several kinds of hybrid topologies are explained in the following sections.

Star-Wired Ring

The **star-wired ring topology** uses the physical layout of a star in conjunction with the token–passing data transmission method. In Figure 5-5, which depicts this architecture, the solid lines represent a physical connection and the dotted lines represent the flow of data. Data are sent around the star in a circular pattern. This hybrid topology benefits from the fault tolerance of the star topology (data transmission does not depend on each workstation to act as a repeater) and the reliability of token passing. Modern Token Ring networks, as specified in IEEE 802.5, use this hybrid topology.

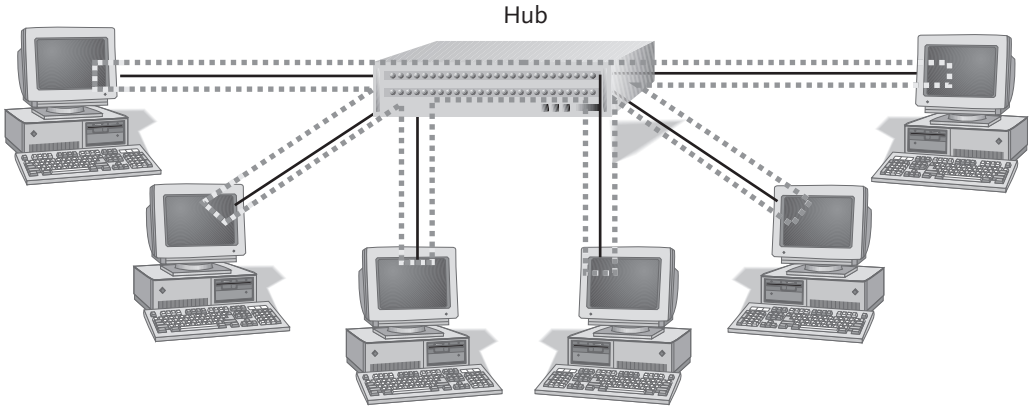

Figure 5-5 A star-wired ring topology network

Star-Wired Bus

Another popular hybrid topology combines the star and bus formations. In a **star-wired bus topology**, groups of workstations are star-connected to hubs and then networked via a single bus, as shown in Figure 5-6. With this design, you can cover longer distances and easily interconnect or isolate different network segments. One drawback is that this option is more expensive than using either the star or, especially, the bus topology alone because it requires more cabling and potentially more connectivity devices. The star-wired bus topology commonly forms the basis for modern Ethernet and Fast Ethernet networks.

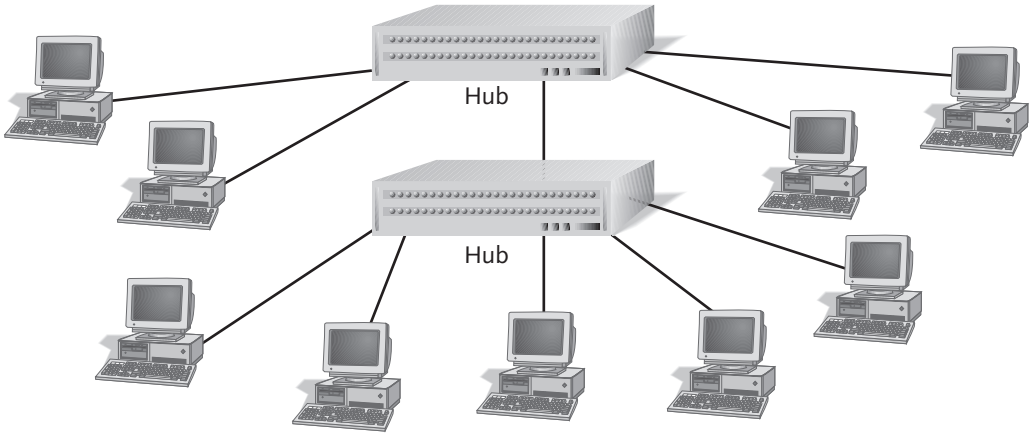

Figure 5-6 A star-wired bus topology network

Daisy-Chained

Even the star-wired ring and bus network topologies are too simplistic to represent a typical medium-sized LAN. Nevertheless, hubs that service star-wired bus or ring topologies can be daisy-chained to form a more complex hybrid topology, as shown in Figure 5-7. A **daisy chain** is a linked series of devices. (As you will learn later in this chapter, in an enterprise network, a daisy-chained network is called a serial backbone network.) Because the star-wired hybrids provide for modular additions, daisy-chaining is a logical solution for growth. Also, because hubs can be easily connected through cables attached to their ports, little additional cost is required to expand a LAN's infrastructure in this way.

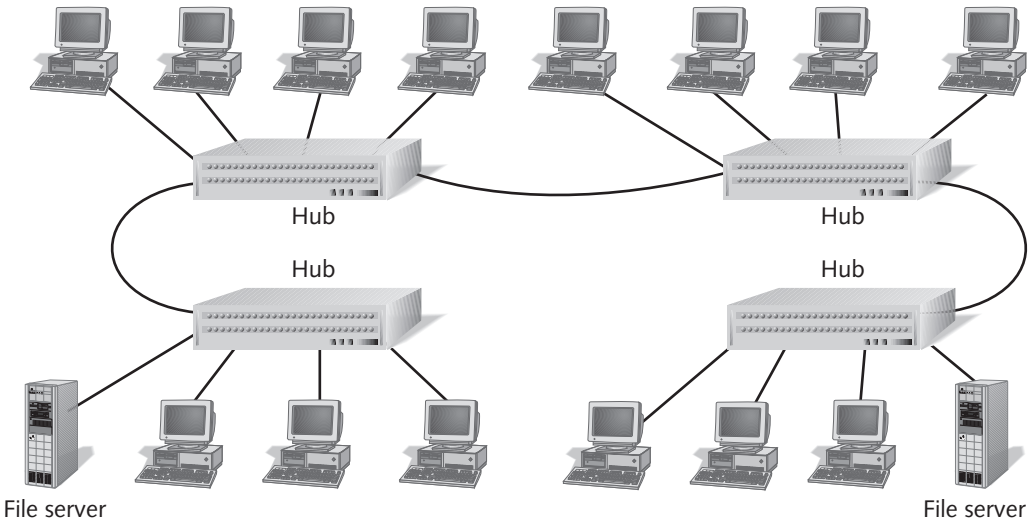

Figure 5-7 Daisy-chained star-wired bus topology

Daisy-chaining simple topologies can present hazards. For example, IEEE specifications such as the Ethernet 802.3 standard dictate the maximum number of hubs that may be connected in sequence to maintain transmission integrity. These standards vary according to the network type. For example, in a 10BaseT network, you may use a maximum of four hubs to connect five network segments. Using more hubs than the standard suggests (in other words, exceeding the maximum network length) will adversely affect the functionality of a LAN. Among other things, if you extend a LAN beyond its recommended size, intermittent and unpredictable data transmission errors will result. Similarly, if you daisy-chain a topology with limited bandwidth, you risk overloading the channel and generating more data errors.

Hierarchical

None of the topologies discussed previously distinguishes between the functions or priorities of the various workgroups. For example, a server that contains a payroll database and serves 50 clients may be attached to the same hub as a workstation that is used only twice each week for data processing. Although both devices are connected to the same hub, they perform vastly different functions. Accordingly, you should assign the payroll server a higher priority for network access. For example, you need to ensure that the payroll server almost never loses network connectivity as a result of a device failure on the network. Thus, you might choose to connect the payroll server directly to the backbone using a more expensive, fault-tolerant hub. The less important, data-processing workstation could be connected to a small, inexpensive hub that is connected to a better hub, which is in turn connected to the backbone. This arrangement minimizes the possibility of the payroll server losing connectivity but increases the possibility of the data-processing workstation losing connectivity.

There are many reasons for separating devices in a hierarchy. You may want to separate hubs, switches, and routers for reasons related to security, cost, scalability, network addressing, bandwidth, or reliability. In addition, there are many ways to separate devices and workgroups, leading to many variations on the hierarchical topology. You can consider the hierarchical topology as similar to an organizational chart in a company, where groups are divided by function, and different personnel belong to different levels in the organizational chart.

One possible way to group devices on a network is to divide them into layers. In the context of topologies, you can think of a layer as a logical division between devices on a network. A hierarchical hybrid topology uses layers to separate devices based on their priority or function. A hierarchical topology may have any number of layers and may connect different types of simple topologies. For example, Figure 5-8 depicts a hierarchical ring topology with three layers. The top layer services the network's backbone, while the second layer provides direct connectivity for the file server ring and intermediate connectivity to the third layer. The third layer then services multiple workgroups, such as Administration, Sales, and IT.

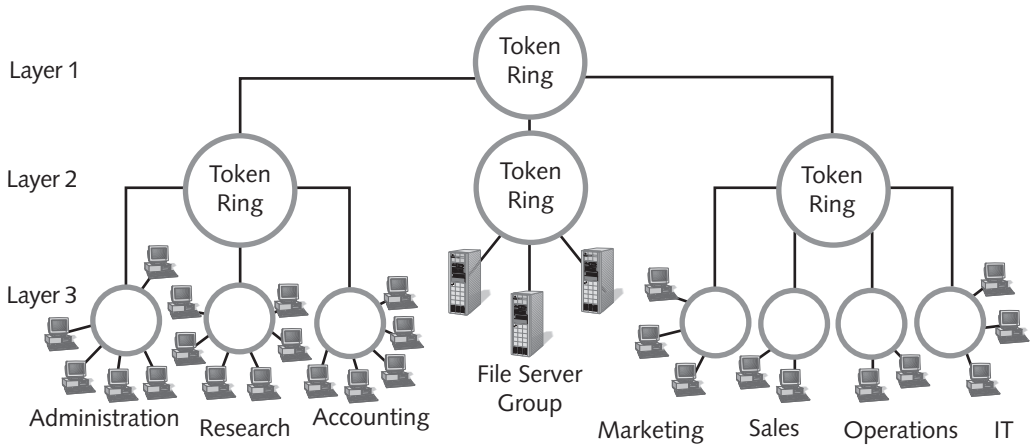

Figure 5-8 Hierarchical ring topology

Arranging topologies in a hierarchy offers several advantages: the ability to segregate bandwidth among different groups, ease in adding or isolating different network groups, and the flexibility to interconnect different network types. For these reasons, hierarchical topologies underlie high-speed LAN and WAN designs.

ENTERPRISE-WIDE TOPOLOGIES

In networking, the term **enterprise** refers to an entire organization, including its local and remote offices, a mixture of computer systems, and a number of departments. Enterprise-wide computing must therefore take into account the breadth and diversity of a large organization's computer needs. Enterprise-wide networks expand on the simple and hybrid LAN topologies. As a result, their topologies require more interconnection devices and more reliable routes than simple LAN topologies can provide. An enterprise-wide network may include or form part of a WAN, but an enterprise-wide network connects only one organization's resources. A WAN (for example, the Internet) may connect resources from many different organizations.

As with LAN topologies, a number of variations on the basic enterprise-wide topologies exist. This section describes some popular methods of arranging these larger networks.

Backbone Networks

As you learned in Chapter 4, a network backbone is the cabling that connects the hubs, switches, and routers on a network. Backbones usually are capable of more throughput than the cabling that connects workstations to hubs. This added capacity is necessary because backbones carry more traffic than any other cabling in the network. For example, an increasing number of businesses are implementing fiber-optic backbone but continue to use CAT5 wiring for the cabling from hubs to workstations. Although even the simplest LAN

(including a star or bus topology LAN) technically has a backbone, enterprise-wide backbones are more complex and more difficult to plan. The backbone is the most significant building block of these networks.

Serial Backbone

A **serial backbone** is the simplest kind of backbone network. It consists of two or more hubs connected to each other by a single cable. Serial backbone networks are identical to the daisy-chained networks discussed in the "Hybrid Physical Topologies" section. As mentioned earlier, they are not suitable for large networks or long distances. Although the serial backbone topology could be used for enterprise-wide networks, it is rarely implemented for that purpose.

Distributed Backbone

A **distributed backbone** consists of a number of hubs connected to a series of central hubs or routers in a hierarchy, as shown in Figure 5-9. In Figure 5-9, the cross-hatched lines represent the backbone. This kind of topology allows for simple expansion and limited capital outlay for growth, because more layers of hubs can be added to existing layers. For example, suppose that you are the network administrator for a small publisher's office. You might begin your network with a distributed backbone consisting of two hubs that supply connectivity to your 20 users, 10 on each hub. When your company hires more staff, you can connect another hub to one of the existing hubs, and use the new hub to connect the new staff to the network.

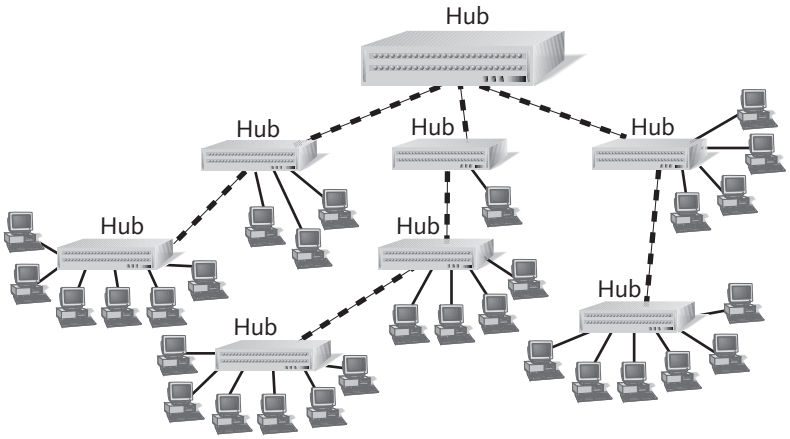

Figure 5-9 A simple distributed backbone network

A more complicated distributed backbone connects multiple LANs or LAN segments using routers, as shown in Figure 5-10. In this example, the routers form the highest layer of the backbone to connect the LANs.

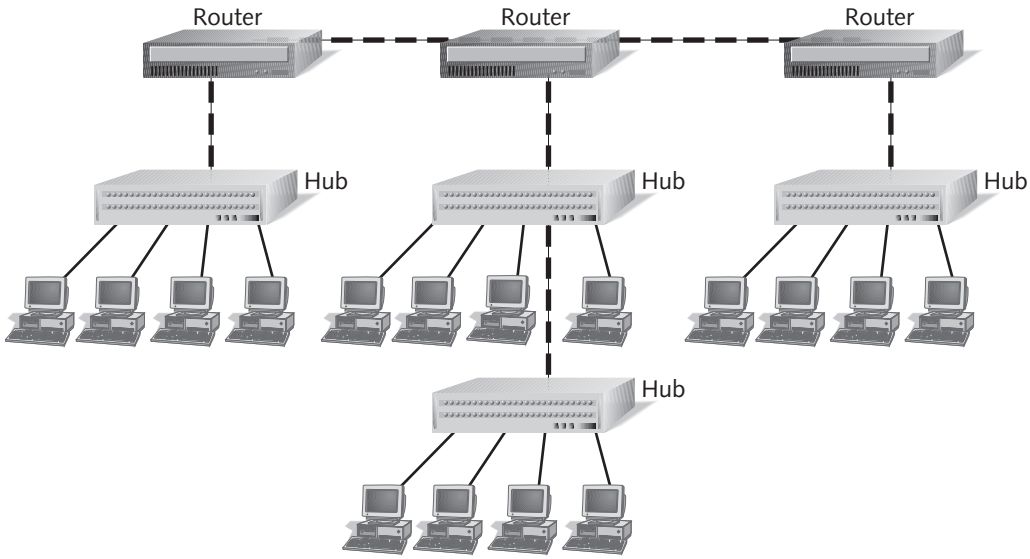

Figure 5-10 A distributed backbone connecting multiple LANs

A distributed backbone also provides network administrators with the ability to segregate workgroups and therefore manage them more easily. It adapts well to an enterprise-wide network confined to a single building, where layers of hubs can be assigned according to the floor or department. When designing a network with a distributed backbone, however, you must consider the maximum allowable distance between nodes and the server dictated by the network media. Another possible problem in this design relates to the central point of failure, the hub at the uppermost layer. Despite these potential drawbacks, implementing a distributed backbone network can be relatively simple, quick, and inexpensive.

Collapsed Backbone

The **collapsed backbone** topology uses a router or switch as the single central connection point for multiple subnetworks, as shown in Figure 5-11. Contrast Figure 5-11 with Figure 5-10, where multiple LANs are connected via a distributed backbone. In a collapsed backbone, a single router or switch is the highest layer of the backbone. The router or switch that makes up the collapsed backbone must contain multiprocessors to handle the heavy traffic going through it. The dangers of using this arrangement relate to the fact that a failure in the central router or switch can bring down the entire network. In addition, because routers cannot move traffic as quickly as hubs, using a router may slow data transmission. (You will learn more about hubs and routers in Chapter 6.)

Nevertheless, a collapsed backbone topology offers substantial advantages. Most significantly, this arrangement allows you to interconnect different types of subnetworks. You can also centrally manage maintenance and troubleshooting chores.

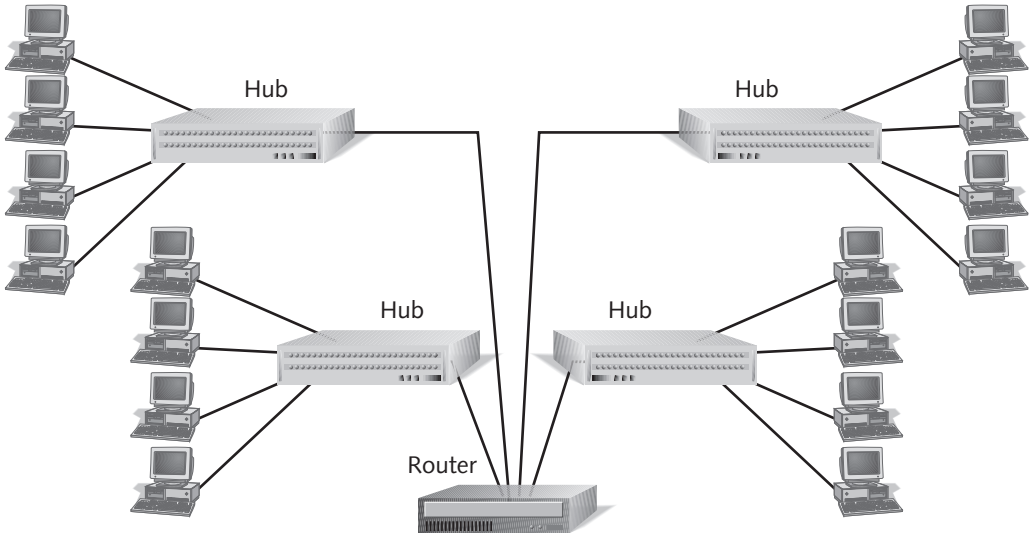

Figure 5-11 A collapsed backbone network

Parallel Backbone

A **parallel backbone** is the most robust enterprise-wide topology. This variation of the collapsed backbone arrangement consists of more than one connection from the central router or switch to each network segment. Figure 5-12 depicts a simple parallel backbone topology. As you can see, each hub is connected to the router or switch by more than one cable. The most significant advantage of using a parallel backbone is that its redundant (duplicate) links ensure network connectivity to any area of the enterprise. Parallel backbones are more expensive than other enterprise-wide topologies because they require much more cabling than the others. However, they make up for the additional cost by offering increased performance.

As a network administrator, you might choose to implement parallel links to only some of the most critical devices on your network. For example, if the first and second hubs in Figure 5-12 connected your Facilities and Payroll departments to the rest of the network, and your organization could never afford to lose connectivity with those departments, you might use a parallel structure for those links. If the third and fourth hubs in Figure 5-12 connected your organization's Recreation and Training departments to the network, you might decide that parallel links were unnecessary for these departments. By selectively implementing the parallel structure, you can lower connectivity costs and leave available additional ports on the connectivity devices.

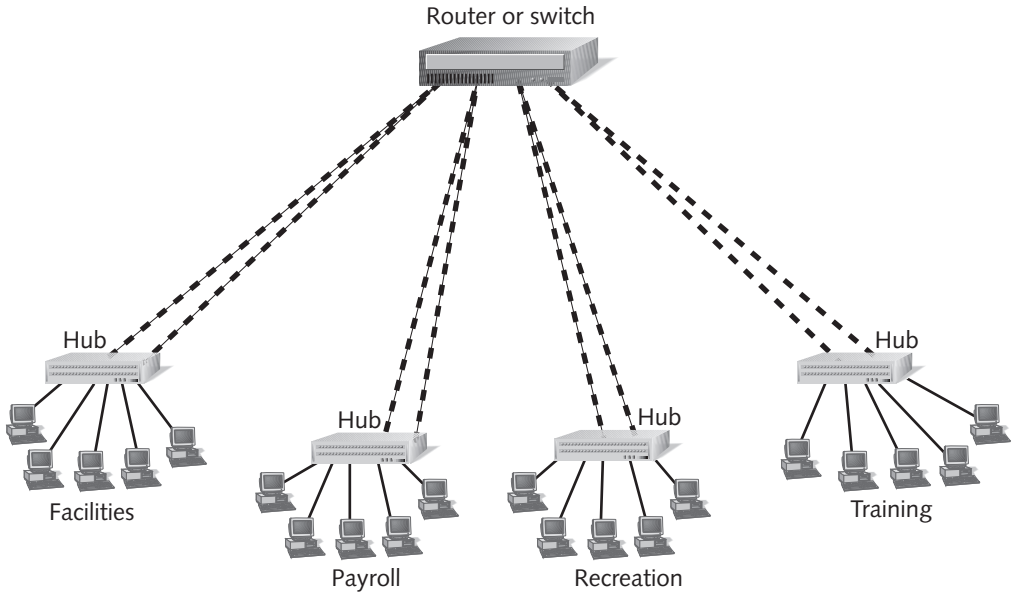

Figure 5-12 A parallel backbone network

Mesh Networks

Of course, backbone networks are not the only type of topology used in enterprise-wide networks. Some organizations use a more intricate topology, known as a mesh network. In a **mesh network**, routers are interconnected with other routers, with at least two pathways connecting each router. (See Figure 5-13.) The mesh network is more complex than the backbone networks. In fact, it typically contains several different backbone networks. Indeed, the term "mesh network" is a general topology term that can apply to many different arrangements of workgroups and interconnection devices.

Although a simple LAN can be a mesh network, most often this topology is employed for enterprise-wide networks and WANs. The Internet is an example of a mesh WAN.

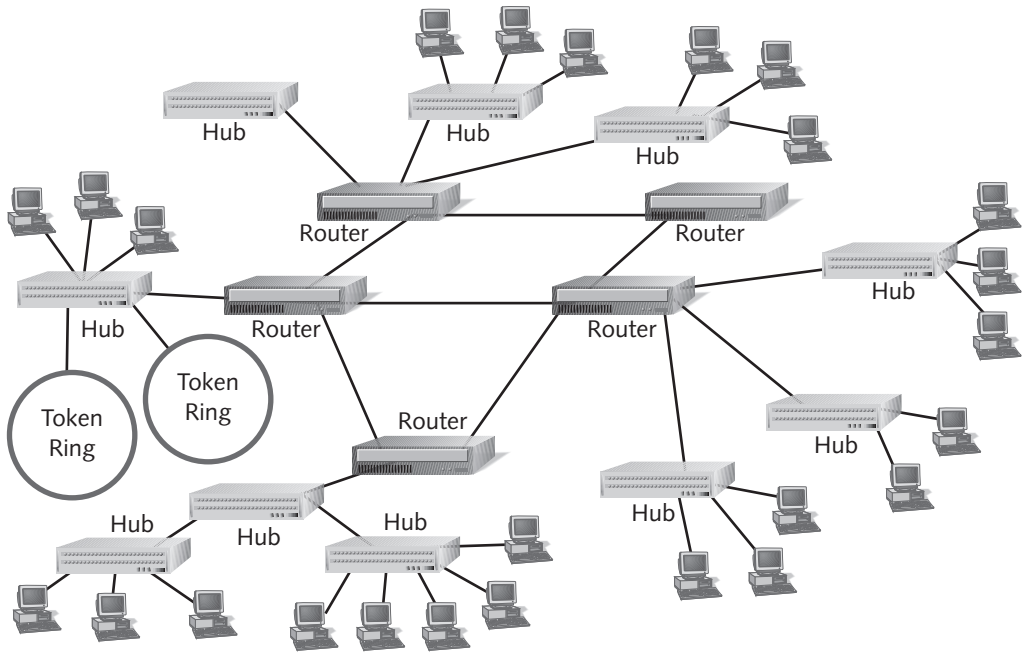

Figure 5-13 An example of a mesh network

WAN Topologies

As you learned in Chapter 1, a wide area network (WAN) is a network connecting geographically distinct locations, which may or may not belong to the same organization. WAN topologies use both LAN and enterprise-wide topologies as building blocks, but add more complexity because of the distance they must cover, the larger number of users they serve, and the heavy traffic they often handle. For example, although a simple ring topology may suffice for a small office with 10 users, it does not scale well and therefore cannot serve 1000 users. The particular WAN topology you choose will depend on the number of sites you must connect, the distance between the sites, and any existing infrastructure.

WAN topologies also differ from LAN topologies in the type and extent of interconnectivity devices they use. This difference is directly related to the fact that networking protocols are handled differently on local network segments than they are on segments involving longer distances. For example, a LAN might carry NetBEUI, IPX/SPX, and TCP/IP traffic over a single segment. Because WANs depend on routers to interconnect LANs, and because NetBEUI is not a routable protocol, however, a WAN link will carry only IPX/SPX and/or TCP/IP traffic. WAN networking technologies, such as ISDN, DSL, and SONET, are discussed in detail in Chapter 7.

Peer-to-Peer

A WAN with single interconnection points for each location is arranged in a **peer-to-peer topology**. A WAN peer-to-peer topology is similar to peer-to-peer communications on a LAN in that each site depends on every other site in the network to transmit and receive its traffic. However, the peer-to-peer LANs use computers with shared access to one cable, whereas the WAN peer-to-peer topology uses different locations, each one connected to another one through (usually) dedicated circuits.

The WAN peer-to-peer topology is often the best option for organizations with only a few sites and the capability to use **dedicated circuits**—that is, continuously available communications channels between two access points that are leased from a telecommunications provider, such as an ISP. You will learn more about dedicated circuits, such as T1 or ISDN connections, in Chapter 7. For now, you simply need to know that dedicated circuits make it possible to transmit data regularly and reliably. Figure 5-14 depicts a peer-to-peer WAN using T1 and ISDN connections.

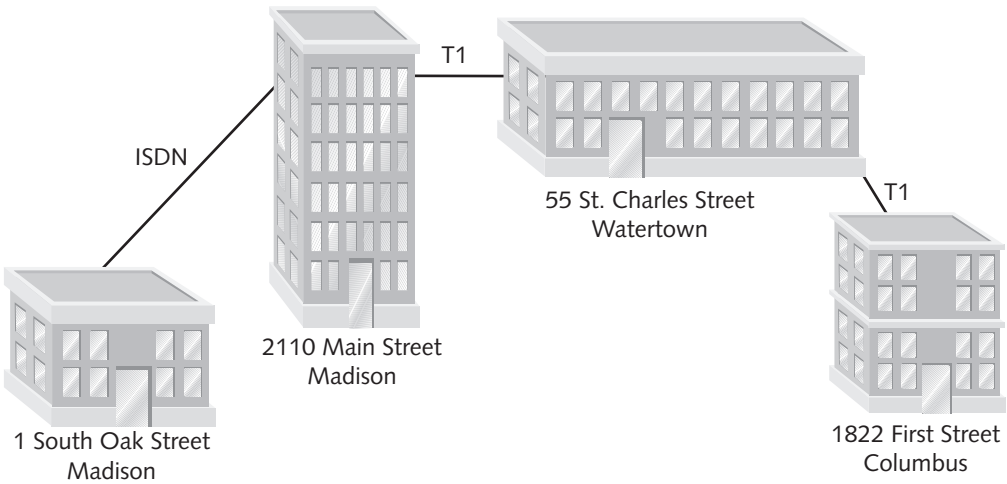

Figure 5-14 A peer-to-peer WAN

Peer-to-peer WAN topologies are suitable for only small WANs. Because all sites must participate in carrying traffic, this model does not scale well. The addition of more sites can cause performance to suffer. Also, a single failure on a peer-to-peer WAN can take down communications between all sites.

Ring

In a **ring WAN topology**, each site is connected to two other sites so that the entire WAN forms a ring pattern, as shown in Figure 5-15. This architecture is similar to the ring LAN topology, except that a ring WAN topology connects locations rather than

local nodes. The advantages of a ring WAN over a peer-to-peer WAN are twofold: a single cable problem will not affect the entire network, and routers at any site can redirect data to another route if one route becomes too busy. On the other hand, expanding ring-configured WANs can be difficult, and it is more expensive than expanding a peer-to-peer WAN because it requires at least one additional link. For these reasons, WANs that use the ring topology are only practical for connecting fewer than four or five locations.

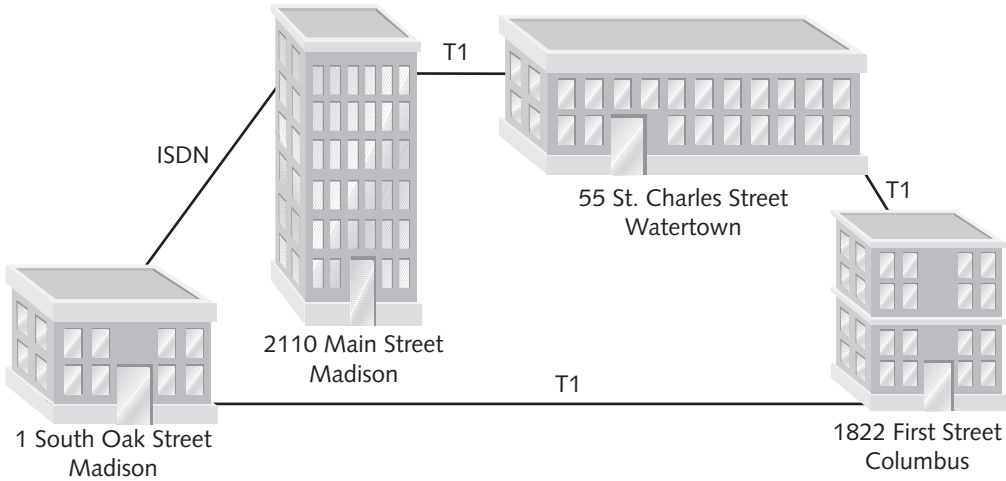

Figure 5-15 A ring-configured WAN

Star

The **star WAN topology** mimics the arrangement of a star LAN. A single site acts as the central connection point for several other points, as shown in Figure 5-16. This arrangement provides separate routes for data between any two sites. As a result, star WANs are more reliable than the peer-to-peer or ring WANs. As a general rule, reliability increases with the number of potential routes data can follow. For example, if the T1 link between the Oak Street and Main Street locations fails, the Watertown and Columbus locations can still communicate with the Main Street location because they use different routes. In a peer-to-peer or ring topology, however, a single failure would halt all traffic between all sites.

Another advantage of a star WAN is that when all of its dedicated circuits are functioning, a star WAN provides shorter data paths between any two sites.

Extending a star WAN is easy, and this expansion costs less than extending a peer-to-peer or ring WAN. For example, if the organization that uses the star WAN pictured in Figure 5-16 wanted to add a Maple Street, Madison, location to its topology, it could simply lease a new dedicated circuit from the Main Street office to its Maple Street office. None of the other offices would be affected by the change. If the organization

were using a peer-to-peer or ring WAN topology, however, two separate dedicated connections would be required to incorporate the new location into the network.

As with star LAN topologies, the greatest drawback of a star WAN is that a failure at the central connection point can bring down the entire WAN. In Figure 5-16, for example, if the Main Street office suffered a catastrophic fire, the entire WAN would fail.

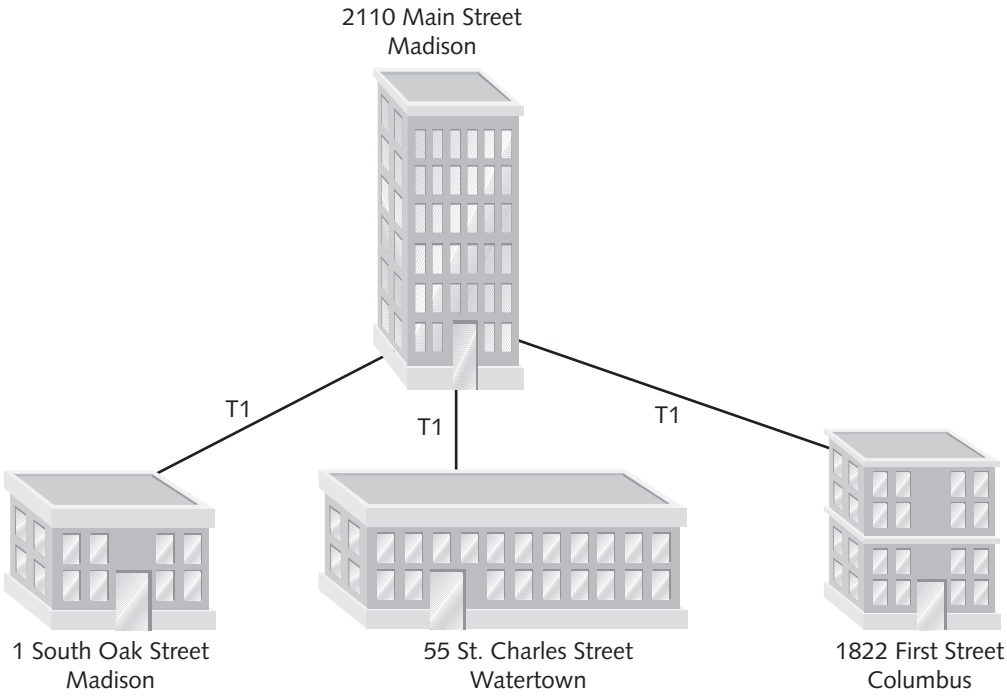

Figure 5-16 A star-configured WAN

Mesh

Like an enterprise-wide mesh, a **mesh WAN topology** incorporates many directly interconnected nodes—in this case, geographical locations. Because every site is interconnected, data can travel directly from its origin to its destination. If one connection suffers a problem, routers can redirect data easily and quickly. Mesh WANs are the most fault-tolerant type of WAN configuration because they provide multiple routes for data to follow between any two points. For example, if the Madison office in Figure 5-17 suffered a catastrophic fire, the Dubuque office could still send and transmit data to and from the Detroit office by going directly to the Detroit office. If both the Madison and Detroit offices failed, the Dubuque and Indianapolis offices could still communicate.

One drawback to a mesh WAN is the cost; connecting every node on a network to every other entails leasing a large number of dedicated circuits. With larger WANs, the expense

can become enormous. To reduce costs, you might choose to implement a partial mesh, in which critical WAN nodes are directly interconnected and secondary nodes are connected through star or ring topologies, as shown in Figure 5-17. Partial-mesh WANs are more practical, and therefore more common in today's business world, than full-mesh WANs.

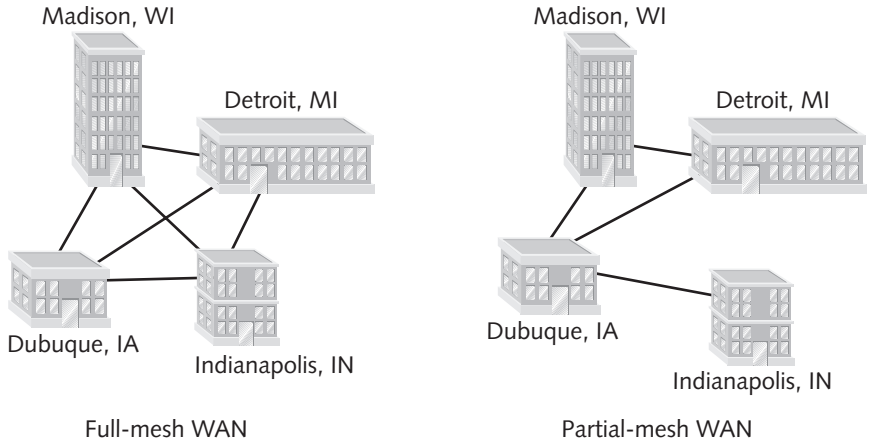

Figure 5-17 Full-mesh and partial-mesh WANs

Tiered

Tiered WAN topologies are similar to the hierarchical hybrid topologies used with LANs. In a **tiered WAN topology**, WAN sites connected in star or ring formations are interconnected at different levels, with the interconnection points being organized into layers. Figure 5-18 depicts a tiered WAN. In this example, the Madison, Detroit, and New York offices form the upper tier, and the Dubuque, Indianapolis, Toronto, Toledo, Washington, and Boston offices form the lower tier. If the Detroit office suffers a failure, the Toronto and Toledo offices cannot communicate with any other nodes on the WAN. Similarly, the Dubuque and Indianapolis offices depend on the Madison office for their WAN connectivity, just as the Washington and Boston offices depend on the New York office for their connectivity.

Variations on this topology abound. Indeed, flexibility makes the tiered approach quite practical. A network architect can determine the best placement of top-level routers based on traffic patterns or critical data paths. In addition, tiered systems allow for easy expansion and inclusion of redundant links to support growth. On the other hand, their enormous flexibility means that creation of tiered WANs requires careful consideration of geography, usage patterns, and growth potential.

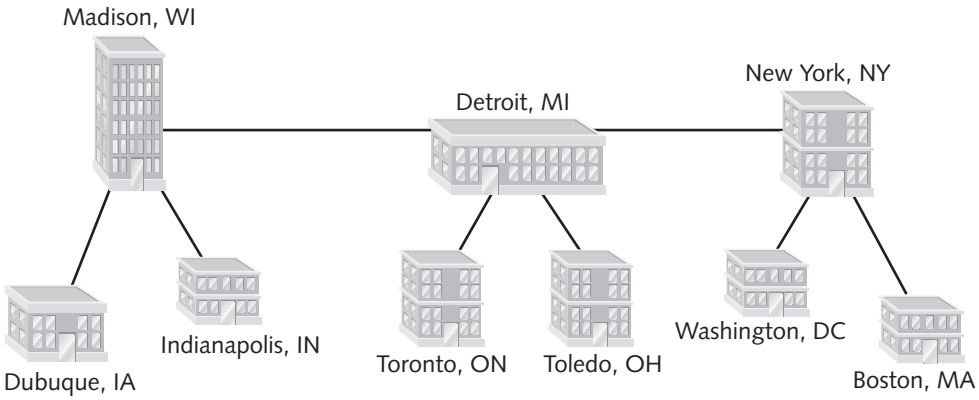

Figure 5-18 A tiered WAN topology

Logical Topologies

The term **logical topology** refers to the way in which data are transmitted between nodes, rather than the layout of the paths that data take. A logical topology may also be called a **network transport system**. The logical topology describes the way data are packaged in frames and the way electrical pulses are sent over the network's physical media. A logical topology's elements reside in both the Data Link and Physical layers of the OSI Model. In contrast, a network's physical topology (which consists of a network's cabling) does not belong to any OSI Model layer, but provides a foundation for the Physical layer. Logical and physical topologies are related, however. Each logical topology not only has its own set of data signaling principles, but also imposes unique requirements on the networking media and physical topology.

The two most popular, traditional network transport systems, or logical topologies, are Ethernet and Token Ring. In Chapter 2, while learning how data are transformed through the OSI Model layers, you were introduced to Ethernet and Token Ring technologies. This section explains both Ethernet and Token Ring systems in more detail. Other logical topologies include LocalTalk, FDDI, and ATM, all of which are also discussed in this chapter.

Switching

Before you can understand specific network transport systems, such as Ethernet and Token Ring, you must be familiar with the concept of switching. **Switching** is a component of a network's logical topology that determines how connections are created between nodes. You will learn more about switches, the hardware that manages network switching, in Chapter 6. For now, you should be aware of the three methods for switching: circuit switching, message switching, and packet switching. Every network transport system relies on one of these switching mechanisms.

Circuit Switching

In **circuit switching**, a connection is established between two network nodes before they begin transmitting data. Bandwidth is dedicated to this connection and remains available until the users terminate communication between the two nodes. While the nodes remain connected, all data follow the same path initially selected by the switch. When you place a telephone call, for example, your call goes through a circuit-switched connection.

Because circuit switching monopolizes its piece of bandwidth while the two stations remain connected (even when no actual communication is taking place), it is not an economical technology. Some network applications that cannot tolerate the time delay it takes to reorganize data packets, such as live audio or videoconferencing, benefit from such a dedicated path, however. When you connect your home PC via modem to your Internet service provider's access server, that connection uses circuit switching. WAN technologies, such as ISDN and T1 service, which are discussed in Chapter 7, also use circuit switching. Finally, ATM, which is discussed later in this chapter, uses circuit switching as well.

Message Switching

Message switching establishes a connection between two devices, transfers the information to the second device, and then breaks the connection. The information is stored and forwarded from the second device once a connection between that device and a third device on the path is established. This "store and forward" routine continues until the message reaches its destination. All information follows the same physical path; unlike with circuit switching, however, the connection is not continuously maintained. E-mail systems use message switching. Message switching requires that each device in the data's path have sufficient memory and processing power to accept and store the information before passing it to the next node. None of the network transmission technologies discussed in this chapter uses message switching.

Packet Switching

A third method for connecting nodes on a network is packet switching. **Packet switching** breaks data into packets before they are transported. Packets can travel any path on the network to their destinations, because each packet contains the destination address and sequencing information. Consequently, packets can attempt to find the fastest circuit available at any instant. They need not follow each other along the same path, nor must they arrive at their destination in the same sequence as when they left the transmitting node.

To understand this technology, imagine that you organized a field trip for 50 colleagues to the National Air and Space Museum in Washington, DC. You gave the museum's precise address to your colleagues and told them to leave precisely at 7:00 A.M. from your office building across town. You did not tell your coworkers which route to take. Some

might choose the subway, others might hail a taxicab, and still others might choose to drive their own cars. All of them will attempt to find the fastest route to the museum. But if a group of six decide to take a taxicab and only four people fit in that taxi, the next two people have to wait for a taxi. Or a taxi might get caught in rush hour traffic and be forced to find an alternate route. Thus, the fastest route might not be evident upon departure. But no matter which transportation method your colleagues choose, you will all arrive at the museum and reassemble as a group. This analogy illustrates how packets travel in a packet-switched network.

The destination node on a packet-switched network reassembles the packets based on their control information. Because of the time it takes to reassemble the packets into a message, packet switching is not suited to live audio or video transmission. Nevertheless, it is a fast and efficient mechanism for transporting typical network data, such as word-processing or spreadsheet files. The greatest advantage to packet switching lies in the fact that it does not waste bandwidth by holding a connection open until a message reaches its destination, as circuit switching does. And unlike message switching, it does not require devices in the data's path to process any information. Examples of packet-switched networks include Ethernet and FDDI. The Internet is also an example of a packet-switched network.

Now that you are familiar with the various types of switching, you are ready to investigate specific logical topologies that may make use of switching.

ETHERNET

As you learned in Chapter 2, Ethernet is a logical topology originally developed by Xerox in the 1970s and later improved by Xerox, Digital Equipment Corporation (DEC), and Intel. This flexible technology can run on a variety of network media and offers excellent throughput at a reasonable cost. Ethernet is, by far, the most popular logical topology for LANs today, and its popularity continues to grow.

Ethernet has evolved through many variations, and continues to improve. As a result of this history, it supports many different versions—so many, in fact, that you will probably find the many variations a little confusing. In Chapter 4 you learned how Ethernet networks may differ at the Physical layer. In this section you will learn how Ethernet networks may differ at the Data Link layer. First, however, you will learn about CSMA/CD, the network access method that all Ethernet networks have in common.

Carrier Sense Multiple Access with Collision Detection (CSMA/CD)

The defining characteristic of Ethernet is its **access method**, or its method of controlling how network nodes access the communications channel. The access method used in Ethernet is called **Carrier Sense Multiple Access with Collision Detection (CSMA/CD)**. All Ethernet networks, independent of their speed or frame type, rely on CSMA/CD. To understand Ethernet, you must first understand CSMA/CD.

Take a minute to think about the full name "Carrier Sense Multiple Access with Collision Detection." The term "Carrier Sense" refers to the fact that Ethernet NICs listen on the network and wait until they detect (or sense) that no other nodes are transmitting data over the signal (or carrier) on the communications channel before they begin to transmit. The term "Multiple Access" refers to the fact that several Ethernet nodes can be connected to a network and can monitor traffic, or access the media, simultaneously.

In CSMA/CD, when a node wants to transmit data it must first access the transmission media and determine whether the channel is free. If the channel is not free, it waits and checks again after a random (but very brief) amount of time. If the channel is free, the node transmits its data. Any node can transmit data once it determines that the channel is free. But what if two nodes simultaneously check the channel, determine that it's free, and begin to transmit? When this happens, their two transmissions will interfere with each other; this is known as a **collision**. In this event, the network performs a series of steps known as the collision detection routine. If a station's NIC determines that its data has been involved in a collision, it will immediately stop transmitting. Next, in a process called **jamming**, it will issue a special 32-bit sequence that indicates to the rest of the network nodes that the station's previous transmission was faulty and that they should not accept those data frames as valid. After waiting, the node will determine if the line is again available; if it is available, the line will retransmit its data.

On heavily trafficked networks, collisions are fairly common. Not surprisingly, the more nodes transmitting data on a network, the more collisions will take place (although a collision rate greater than 5% of all traffic is unusual and may point to a problematic NIC or poor cabling on the network). When an Ethernet network grows to include a particularly large number of nodes, you may see performance suffer as a result of collisions. This "critical mass" number depends on the type and volume of data that the network regularly transmits. Collisions can corrupt data or truncate data frames, so it is important that the network detect and compensate for them. Figure 5-19 depicts the way CSMA/CD regulates data flow to avoid and, if necessary, detect collisions.

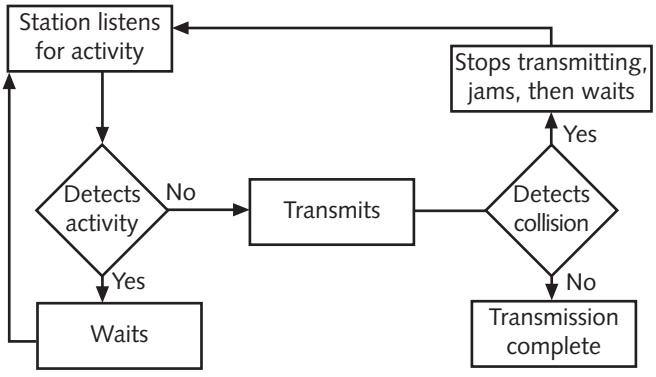

Figure 5-19 CSMA/CD process

On an Ethernet network, an individual segment is known as a **collision domain**, or a portion of a network in which collisions will occur if two nodes transmit data at the same time. When designing an Ethernet network, it's important to note that, since repeaters simply regenerate any signal they receive, they repeat collisions just as they repeat data. Thus, connecting multiple parts of a network with repeaters results in a larger collision domain. Higher-layer connectivity devices, such as switches and routers, however, can separate collision domains.

Collision domains play a role in the Ethernet cabling distance limitations. For example, if the distance between two nodes on a segment connected to the same 100BaseT network bus exceeds 100 meters, data propagation delays will be too long for CSMA/CD to work. A **data propagation delay** is the length of time data take to travel from one point on the segment to another point. When data takes a long time, CSMA/CD's collision detection routine cannot identify collisions accurately. In other words, one node on the segment might begin its CSMA/CD routine and determine that the channel is free even though a second node has begun transmitting, because the second node's data is taking so long to reach the first node.

In Fast Ethernet, data travel so quickly that NICs can't always keep up with the collision detection and retransmission routines. Because of the speed employed on a 100BaseT network, the window of time for the NIC to both detect and compensate for the error is much less than that of a 10BaseT, 10Base2, or 10Base5 network. To minimize undetected collisions, 100BaseT buses can support a maximum of three network segments connected with two hubs, while 10BaseT buses can support a maximum of five network segments connected with four hubs. This shorter path reduces the highest potential propagation delay between nodes.

Demand Priority

In Chapter 4 you learned about multiple Physical layer specifications for Ethernet networks, including 10BaseT and 100BaseT. All these types rely on CSMA/CD for data flow control. You also learned about a networking technology that is similar to Ethernet—100BaseVG (or 100BaseVG-AnyLAN). One significant difference between 100BaseVG and 100BaseT (and the characteristic that makes it not quite "Ethernet") is that 100BaseVG does not use CSMA/CD, but rather an access method called demand priority. In **demand priority**, each device on a star or hierarchical network sends a request to transmit to the central hub. The hub then grants the requests one at a time. It examines incoming data packets, determines the location of the destination node, and forwards the packets to that destination. Because demand priority runs on a star topology, in which each node is linked directly to a connectivity device, no workstations except the source and destination can "see" the data. Data travel from one device to the hub, then to another device. The hub acts as a central transfer point. Figure 5-20 contrasts CSMA/CD with demand priority techniques.

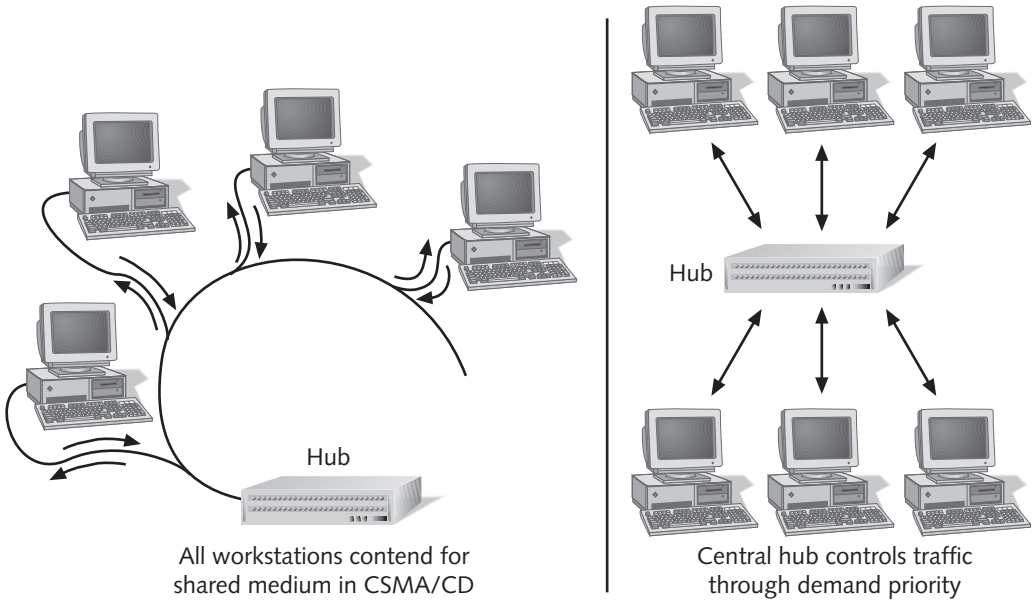

Figure 5-20 CSMA/CD versus demand priority

Because the hub determines which nodes transmit and when, data collisions do not happen on demand priority networks. Data travel unimpeded, with no need for the collision detection and correction required by CSMA/CD. Another advantage to demand priority networks is that, because data do not pass by each node on the network, they remain secure; packets transmitted from one workstation cannot be trapped and decoded by any workstation on the network except for the destination. In addition, in a demand priority network, the hub can prioritize transmission requests. If multiple requests arrive at the hub simultaneously, the hub services the highest-priority request first. This approach allows 100BaseVG to better serve networks that carry audio, video, or other time-sensitive data. (The "VG" part of its name stands for voice grade specification.) Before you can use demand priority on your network, you must ensure that NIC drivers compatible with the 100BaseVG priority assignment scheme have been installed.

Demand priority requires an **intelligent hub**—that is, a hub that can manage transmissions by dictating which nodes send and receive data at every instant—rather than a hub that simply regenerates signals. Some Ethernet networks do not have intelligent hubs. Another disadvantage of a demand priority network is that the time the hub takes to process each request reduces the network's overall performance, so that a 100BaseVG network usually cannot match the speed of a 100BaseT (CSMA/CD) network. And as you learned in Chapter 4, 100BaseVG networks cannot take advantage of full duplexing, which can potentially double a network's bandwidth. For these reasons (and because compatible hardware is difficult to find), 100BaseVG networks based on demand priority are uncommon.

Switched Ethernet

Traditional Ethernet LANs, called **shared Ethernet**, supply a fixed amount of bandwidth that must be shared by all devices on a segment. Stations cannot send and receive data simultaneously, nor can they transmit a signal when another station on the same segment is sending or receiving data. This is because they share a segment and a hub or repeater, which merely amplifies and retransmits a signal over the segment. In contrast, a **switch** is a device that can separate a network segment into smaller segments, with each segment being independent of the others and supporting its own traffic. **Switched Ethernet** is a newer Ethernet model that enables multiple nodes to simultaneously transmit and receive data over different logical network segments. By doing so, each node can individually take advantage of more bandwidth. Figure 5-21 shows how switches can isolate network segments. You will learn more about switches in Chapter 6.

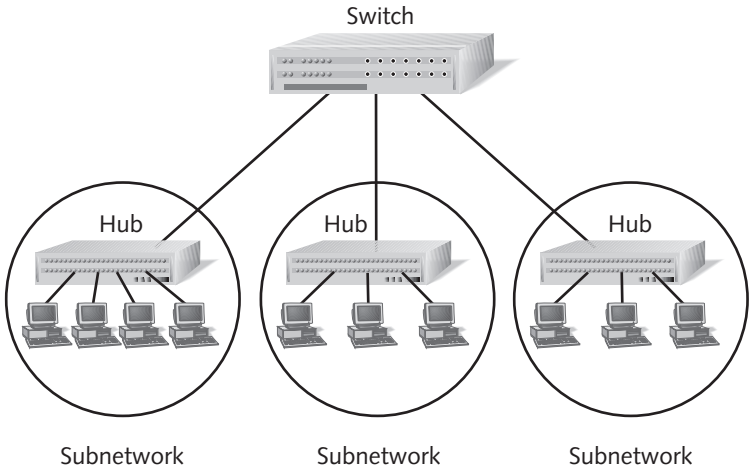

Figure 5-21 A switched Ethernet network

Using switched Ethernet increases the effective bandwidth of a network segment because fewer workstations must vie for the same time on the wire. In fact, applying switches to a 10-Mbps Ethernet LAN can increase its effective data transmission rate to 100 Mbps. For organizations with existing 10BaseT infrastructure, switches offer a relatively simple and inexpensive way to add bandwidth. Switches can be placed strategically on an organization's network to balance traffic loads and reduce congestion.

Note, however, that switches are not always the best answer to heavy traffic and a need for greater speeds. In a case where an enterprise-wise Ethernet LAN is generally overtaxed, you should consider upgrading the network's design or infrastructure.

Gigabit Ethernet

As you would probably guess, the evolution of Ethernet has not stopped with the development of switched Ethernet and the 100-Mbps standard. IEEE established specifications

for an Ethernet version that runs at 1000 Mbps, called **1 Gigabit Ethernet**. 1 Gigabit Ethernet can technically run over unshielded twisted-pair (UTP) cable. However, it performs much better over (multimode) fiber, which is specified by the IEEE 802.3z project. A segment of 1 Gigabit Ethernet running on UTP can span a maximum of 100 meters, while a segment running on fiber can span a maximum of 550 meters. Like Fast Ethernet, a fiber-based 1 Gigabit Ethernet network uses CSMA/CD transmission and the IEEE 802.3 frame type (discussed below) and is capable of full duplexing.

You will most likely encounter 1 Gigabit Ethernet as part of a network's backbone. It is well suited to connecting multiple buildings on a single campus, for example. Currently, this scheme would not be appropriate for connecting workstations to hubs, for example, because workstations' NICs and CPUs could not process data fast enough to make the cost worthwhile. In the near future, however, PCs will be equipped with adequate hardware and processing power to take advantage of 1 Gigabit Ethernet.

But the race for greater throughput has not stopped at 1 gigabit. In March 1999 representatives from the networking industry began discussing a **10 Gigabit Ethernet** standard. The standards for 10 gigabits are currently being defined by the IEEE 802.3ae committee and will include full-duplexing and multimode fiber requirements. IEEE is aiming to make the 10 Gigabit standard compatible with the Physical layer standards for 1 Gigabit Ethernet to allow organizations to easily upgrade their networks. The 1- and 10-Gigabit technologies will compete directly with other fast networking solutions such as Asynchronous Transfer Mode (ATM), which is covered later in this chapter.

Ethernet Frame Types

Chapter 2 introduced you to data frames, the packages that carry higher-layer data and control information that enable data to reach their destinations without errors and in the correct sequence. The Ethernet data frame discussed in Chapter 2 is an example of a typical Ethernet data frame. In fact, networks may use one (or a combination) of four kinds of Ethernet data frames: Ethernet IEEE 802.3, Novell Proprietary 802.3, Ethernet II, and IEEE 802.3 SNAP. This variety of Ethernet "types" came about as different organizations released and revised Ethernet standards during the 1980s, changing as LAN technology evolved.

All Ethernet networks, no matter what their frame type, are standardized by the IEEE 802.3 committee and therefore fall under the "802.3" standards. In other words, even if one company calls one particular Ethernet frame type "802.2," the frame falls under the 802.3 networking standard.

Each frame type differs slightly in the way it codes and decodes packets of data traveling from one device to another. The routines that manage Data Link layer functions on a node must be configured to expect one type of frame. On a workstation or server, you can specify the network's frame type through the operating system, usually within the NIC configuration interface. If a node receives a different type of frame than it expects,

it will not be able to decode the data contained in the frame. The result is that a device cannot log onto the network or exchange data. For this reason, it is important for LAN administrators to standardize the type of frame used on their Ethernet networks. This is simple to achieve on networks that use only one type of network operating system and client software, but networks commonly use multiple versions of both.

Ethernet frame types do not depend on the Ethernet's Physical layer specification. In other words, frame types have no relation to the topology or cabling characteristics of the network. Thus, Physical layer standards, such as 10Base2 and 10BaseT, have no effect on the type of framing that occurs in the Data Link layer. At the same time, framing also takes place independent of the higher-level layers. Thus, most types of frames can carry any one of many higher-layer protocols. For example, a single Ethernet IEEE 802.3 data frame may carry either TCP/IP or IPX/SPX traffic (but not both simultaneously).

Ethernet frame sizes vary. However, no matter what its type, each frame contains a 14-byte header and a 4-byte Frame Check Sequence field. These two fields add 18 bytes to the frame size. The data portion of the frame may contain from 46 to 1500 bytes of information. (If less than 46 bytes of data are carried, the network fills out the data portion with extra bytes until it totals 46 bytes. The extra bytes are known as **padding** and have no effect on the data being transmitted.) Thus the minimum Ethernet frame size is 18 + 46, or 64, bytes and the maximum Ethernet frame size is 18 + 1500, or 1518, bytes. Because of the overhead present in each frame and the time required to enact CSMA/CD, the use of larger frame sizes on a network generally results in faster throughput. To some extent, you cannot control frame sizes. You can, however, minimize the number of broadcast frames on your network, which is desirable because broadcast frames tend to be very small and, therefore, inefficient.

Each Ethernet frame also contains a 7-byte preamble. The original Ethernet frame type (which predated the IEEE standard) calls for an 8-byte preamble. The IEEE 802.3 standard changed the last byte of this preamble to a start-of-frame delimiter (SFD), which identifies where the data field begins. Thus, the preamble itself is only 7 bytes, but taken with the SFD, it is equivalent to the original 8-byte preamble. Now that you understand the similarities between different Ethernet frame types, you can learn about what distinguishes each type.

IEEE 802.3 ("Ethernet 802.2," or "LLC")

IEEE 802.3 frame is the default frame type for versions 4.x and higher of the Novell NetWare network operating system. It is the most popular Ethernet frame type for use with IPX/SPX traffic on most contemporary LANs. The defining characteristics of its data portion are the source and destination service access points that belong to the Logical Link Control (LLC) layer, a sublayer of the Data Link layer. Because the IEEE 802.3 Ethernet frame includes this LLC element, it is sometimes called an **LLC frame**. In Novell's lexicon, which was adopted by many other companies, this frame is called an **Ethernet 802.2 frame**. Figure 5-22 depicts an IEEE 802.3 frame.

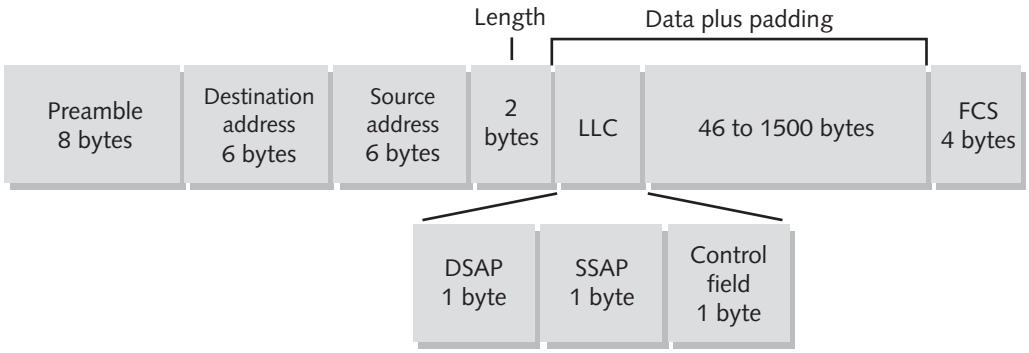

Figure 5-22 An IEEE 802.3 frame

Like other Ethernet frame types, an IEEE 802.3 frame contains an 8-byte preamble. This preamble signals the receiving node that data are incoming and indicates when the data flow is about to begin. Preambles are not included when you calculate a frame's total size.

The destination address and source address fields in an IEEE 802.3 frame are each 6 bytes long. As you might guess, the destination address identifies the recipient of the data frame, and the source address identifies the network node that originally sent the data. Recall from Chapter 3 that any network device can be identified by its logical address (protocol-dependent) or its physical address (hardware-dependent). The physical address is also called the Medium Access Control (MAC) address. Because MAC addresses are hard-coded into the node's NIC, and each manufacturer uses a different identifying code, no two devices should ever have the same address. The source address and destination address fields of an IEEE 802.3 frame use the MAC address to identify where data originated and where it should be delivered. The same is true for all Ethernet frame types.

IEEE 802.3 frames also include a field 2 bytes long that identifies the length of the data field. The data field in an IEEE 802.3 frame contains not only the data transmitted by the source node, but also Logical Link Control (LLC) layer information whose purpose is to distinguish among multiple clients on a network. It may also include padding, if the LLC and data information do not total at least 46 bytes. The length field, however, does account for padding. It will report only the length of LLC plus data information.

The LLC information comprises three fields: Destination Service Access Point (DSAP), Source Service Access Point (SSAP), and a control field. Each of these fields is 1 byte long, making the total LLC field 3 bytes long. A **Service Access Point (SAP)** identifies a node or internal process that uses the LLC protocol. Each process between a source and destination node on the network may have a unique SAP. The control field identifies the kind of LLC connection that must be established, from unacknowledged (connectionless) to fully acknowledged (connection-oriented).

The data field of the IEEE 802.3 frame is the easiest field to understand. It contains the data sent by the originating node, before the packet was passed down from the top layer of the OSI Model.

The **Frame Check Sequence (FCS)** field ensures that the data are received just as they were sent. When the source node transmits the data, it performs an algorithm (a mathematical routine) called a **Cyclical Redundancy Check (CRC)**. CRC takes the values of all of the preceding fields in the frame and generates a unique 4-byte number, the FCS. When the destination node receives the frame, it unscrambles the FCS via CRC and makes sure that the frame's fields match their original form. If this comparison fails, the receiving node assumes that the frame has been damaged in transit and requests that the source node retransmit the data.

Novell Proprietary 802.3 (or "Ethernet 802.3")

The **Novell proprietary 802.3 frame** type is the original NetWare frame type and the default frame type for networks running NetWare versions lower than 3.12. It supports only the IPX/SPX protocol. The Novell proprietary 802.3 frame type is sometimes also called **802.3 Raw** because its data portion contains no control bits. Novell and other companies also call this type of Ethernet frame simply **Ethernet 802.3 frame**. Its fields match those of IEEE 802.3, minus the Logical Link Control layer information. Figure 5-23 depicts a Novell proprietary 802.3 frame.

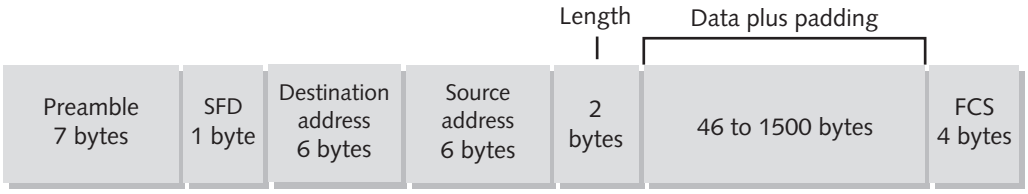

Figure 5-23 A Novell proprietary 802.3 frame

Since the release of later versions of NetWare, most organizations that rely on Novell LANs have migrated their frame types to IEEE 802.3. Thus, the Novell proprietary 802.3 frame type is rarely used on modern networks.

Ethernet II

Ethernet II frame was the original Ethernet frame type developed by DEC, Intel, and Xerox, before the IEEE began to standardize Ethernet. The Ethernet II frame is similar to the Novell proprietary 802.3 frame, in that it lacks LLC information. Ethernet II frames contain a 2-byte type field, however, whereas the IEEE 802.3 and the Novell proprietary 802.3 frames contain a 2-byte length field. This type field identifies the upper-layer protocol contained in the frame. For example, IPX uses a type field of 8137. IP uses a type field of 0800. This field enables Ethernet II to support Novell IPX/SPX, TCP/IP, and AppleTalk protocols, and it compensates for the lack of LLC information. Figure 5-24 depicts an Ethernet II frame.

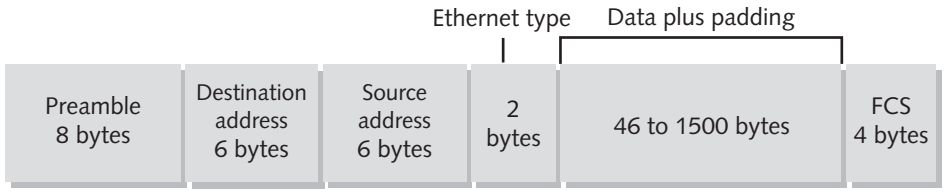

Figure 5-24 An Ethernet II frame

IEEE 802.3 SNAP

IEEE 802.3 SNAP frame is an adaptation of IEEE 802.3 and Ethernet II. "SNAP" stands for Sub-Network Access Protocol. The SNAP portion of the frame is what IEEE 802.3 SNAP borrowed from IEEE 802.3—the three LLC fields (DSAP, SSAP, and the Control field). The IEEE 802.3 SNAP frame, however, contains an additional field: the Organization ID (OUI), a method of identifying the type of network on which the frame is running. In addition, IEEE 802.3 SNAP frames carry Ethernet type information, just as an Ethernet II frame does. IEEE 802.3 SNAP is compatible with IPX/SPX, TCP/IP, and AppleTalk protocols, but it is rarely used on contemporary LANs. Figure 5-25 depicts an IEEE 802.3 SNAP frame.

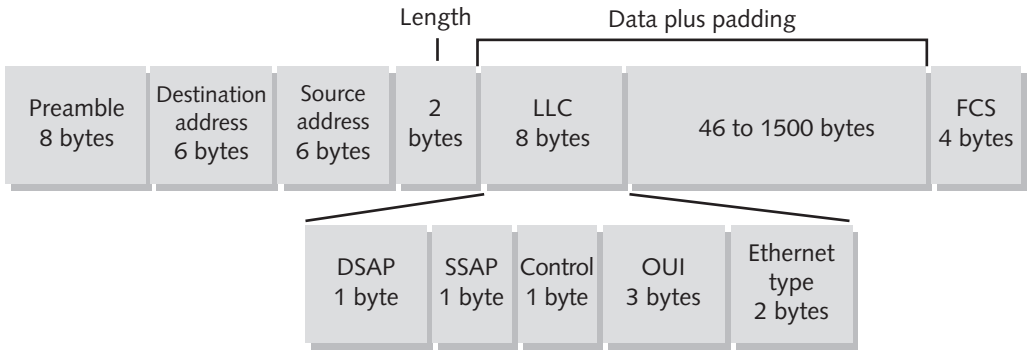

Figure 5-25 An IEEE 802.3 SNAP frame

Understanding Frame Types

You might wonder why you should learn about frame types, which represent the underlying structure of data signals. It's a good question, and the answer relates to problem solving in networks. As a networking professional, you may need to capture data and analyze frames when troubleshooting with a protocol analyzer. When analyzing the frames, you can actually decode the data in every field, if you know what the fields are.

Learning about networks is analogous to learning a foreign language, with the frame type being the language's syntax. Just as you might know the Japanese word for "go," but not know how to use it in a sentence, you may know all about the IPX/SPX protocol, but not how devices handle it. (Chapter 12 covers this kind of troubleshooting in more detail.)

A good knowledge of frame types will help you in many areas of your networking career. For example, to improve a network's performance, you might need to identify the kinds of frames that traverse the network. When working with switches and routers, you might have to configure the device to handle a certain frame type. Or you might decide to configure a switch to accept all types of frames, with the understanding that network performance will suffer as a result of the device having to examine each incoming packet. Probably the most common problem relating to frame types arises from incompatibility between the frame type a workstation expects to receive and the frame type the server actually transmits.

As you learned in Chapter 2, you can use multiple frame types on a network, but you cannot expect interoperability between the frame types. For example, in a mixed environment of NetWare 3.11 and NetWare 4.11 servers, your network will probably support both Ethernet 802.3 and Ethernet 802.2 frames. A workstation connecting to the NetWare 3.11 server might be configured to use the Ethernet 802.3 frame, while a workstation connecting to the NetWare 4.11 server might use Ethernet 802.2.

Modern networks simplify the issue of frame type specification by allowing you to instruct a NIC, through the device driver software, to automatically sense what types of frames are running on a network and adjust themselves to that specification. This feature, called **autosense**, is generally available on all NICs manufactured in the last few years. Workstations, networked printers, and servers added to an existing network can all take advantage of autosense. Even if you use autosense, you should nevertheless know what frame types are running on your network so that you can troubleshoot connectivity problems. As easy as it is to configure, the autosense feature is not infallible.

Design Considerations for Ethernet Networks

In the previous chapter you learned about the Physical layer characteristics of different Ethernet versions. For reference, the following list summarizes essential information from that chapter and from the preceding section about this most important type of network.

- *Cabling*—Ethernet networks can use coaxial cable or unshielded twisted-pair cabling.
- *Connectivity devices*—Ethernet NICs, switches, hubs, routers, and bridges are generally less expensive than comparable Token Ring or LocalTalk equipment.
- *Number of stations*—The number of allowable stations on a 10BaseT or 100BaseTX Ethernet network is limited to 1024.
- *Speed*—Ethernet networks may have a throughput of 10 Mbps, 100 Mbps, 1 Gbps, and soon, 10 Gbps.
- *Scalability*—You can easily expand Ethernet networks by adding connectivity devices on the bus. However, bear in mind each type of network's size limitations.
- *Topology*—10BaseT and 100BaseTX Ethernet networks use a star-wired bus hybrid topology, which is highly fault-tolerant.

LocalTalk

Now that you have learned about Ethernet, a very common logical topology, you will learn about some less common logical topologies, including LocalTalk. **LocalTalk** is a logical topology designed by Apple Computer, Inc. specifically for networking Macintosh computers. It has been included with the Macintosh operating system since 1984, and it provided a simple, cost-effective way of interconnecting Macintosh devices. However, LocalTalk is only capable of 230 Kbps maximum throughput—much less than the 10-Mbps or 100-Mbps throughput of an Ethernet network. Also, LocalTalk is not easily supported by non-Macintosh devices. And, since Macintosh computers are capable of using Ethernet as a network access method, Ethernet is usually preferred over LocalTalk. An instance in which LocalTalk might still be appropriate is for a home network that requires simple configuration and does not require high throughput. Although you may never need to build a LocalTalk network, the essential details of this logical topology are included here in case you must modify or troubleshoot one.

LocalTalk uses a transmission method called **Carrier Sense Multiple Access/Collision Avoidance (CSMA/CA)**. It is similar to the CSMA/CD used in Ethernet networks, except that a node on a LocalTalk network signals its intent to transmit before it actually does so. In this way, collisions and the need for data retransmittals are (mostly) avoided. At the Physical layer, LocalTalk networks require twisted-pair wiring and, in fact, use the same type of cabling used for telephone connections. They may rely on a star or, more often, a bus physical topology. Their maximum segment length is 1000 feet, or approximately 305 meters. Up to 32 nodes may be connected to any single LocalTalk network before data errors begin to occur.

The latest Macintosh computers don't even include a serial port into which a LocalTalk cable could be inserted. But these computers can communicate with LocalTalk devices (such as an older printer) through a special LocalTalk-to-Ethernet adapter.

To connect a Macintosh device to a LocalTalk network, insert one end of a cable into a serial port (for example, the printer or modem port), and then connect this cable to its transceiver (which, in LocalTalk terminology, is called a **teleconnector**). In case the device is being incorporated into a bus topology, this teleconnector contains a resistor to guard against signal bounce. The teleconnector is then connected with a twisted-pair patch cable to the wall jack, which leads to the network's horizontal wiring.

The Macintosh operating system allows devices connected in this manner to share resources in a peer-to-peer fashion without any additional software. A client-server network can be established by introducing a server and AppleShare software and configuring workstations as clients. As with other logical topologies, such as Ethernet, LocalTalk networks support multiple higher-layer protocols. By default, LocalTalk relies on the AppleTalk protocol (discussed in Chapter 3), but it may also support the Macintosh version of TCP/IP called **MacTCP**. In order to support other versions of TCP/IP, LocalTalk requires that these TCP/IP packets be encapsulated by AppleTalk packets.

TOKEN RING

Now that you have learned about LocalTalk and the many forms of Ethernet, you are ready to learn about Token Ring, a less common, but still important logical topology. As you learned in Chapter 2, Token Ring is a network transport system first developed by IBM in the 1980s. In the early 1990s, the Token Ring architecture competed strongly with Ethernet to be the most popular logical topology. Since that time, the economics, speed, and reliability of Ethernet have improved, leaving Token Ring behind. Because IBM developed Token Ring, some IBM-centric IT departments continue to use it. Many other network managers have changed their former Token Ring networks into Ethernet networks.

Token Ring networks are generally more expensive to implement than Ethernet networks. Proponents of the Token Ring technology argue that, although some of its connectivity hardware is more expensive, its reliability results in less downtime and lower network management costs than Ethernet. On a practical level, Token Ring has probably lost the battle for superiority because its developers were slower to develop a high-speed standard. Token Ring networks can run at either 4, 16, or 100 Mbps. The 100-Mbps Token Ring standard, finalized in 1999, is known as **High-Speed Token Ring (HSTR)**. HSTR can use either twisted-pair or fiber-optic cable as its transmission medium. While it is as reliable and efficient as Fast Ethernet, it is less common because of its more costly implementation.

Token Ring networks use the token-passing routine and a star-ring hybrid physical topology. Recall from the discussion of the ring topology earlier in this chapter that a token designates which station on the ring can transmit information on the wire. On a Token Ring network, one workstation, called the active monitor, acts as the controller for token passing. Specifically, the **active monitor** maintains the timing for ring passing, monitors token and frame transmission, detects lost tokens, and corrects errors when a timing error or other disruption occurs. Only one workstation on the ring can act as the active monitor at any given time.

In token passing, a 3-byte token circulates around the network. When a station has something to send, it picks up the token, changes it to a frame, and then adds the header, information, and trailer fields. The header includes the address of the destination node. All nodes read the frame as it traverses the ring to determine whether they are the intended recipient of the message. If they are, they pick up the data, then retransmit the frame to the next station on the ring. When the frame finally reaches the originating station, the originating workstation reissues a free token that can then be used by another station. The token passing control scheme ensures high data reliability (no collisions) and an efficient use of bandwidth. It also does not impose distance limitations on the length of a LAN segment, unlike CSMA/CD. On the other hand, token ring passing generates extra network traffic.

The Token Ring architecture is often mistakenly described as a pure ring topology. In fact, it uses a star-ring hybrid topology in which data circulate in a ring fashion, but the physical layout of the network is a star.

IEEE standard 802.5 describes the specifications for Token Ring technology. Token Ring networks transmit data at either 4 Mbps, 16, or 100 Mbps over shielded or unshielded twisted-pair wiring. You may have as many as 255 addressable stations on a Token Ring network that uses shielded twisted-pair or as many as 72 addressable stations on one that uses unshielded twisted-pair. All Token Ring connections rely on a NIC that taps into the network through a **Multistation Access Unit (MAU)**, Token Ring's equivalent of a hub. NICs can be designed and configured to run specifically on 4-, 16-, or 100-Mbps networks or they can be designed to accommodate both data transmission rates. In the star-ring hybrid topology, the MAU completes the ring internally with Ring In and Ring Out ports at either end of the unit. In addition, MAUs typically provide eight ports for workstation connections. You can easily expand a Token Ring network by connecting multiple MAUs through by their Ring In and Ring Out ports, as shown in Figure 5-26. Unused ports on a MAU, including Ring In and Ring Out ports, have self-shorting data connectors that internally close the loop.

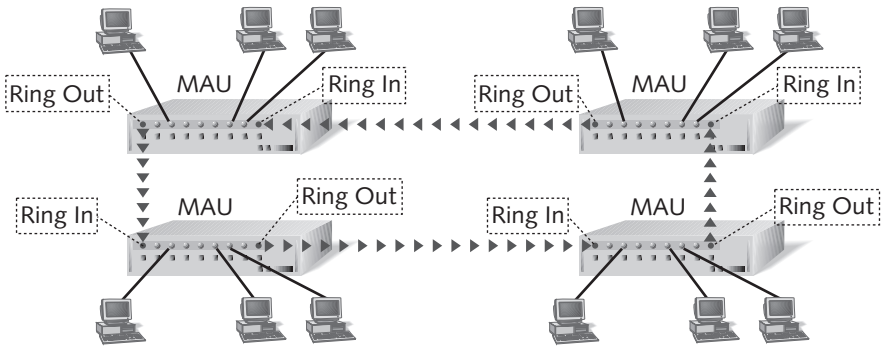

Figure 5-26 Interconnected Token Ring MAUs

The self-shorting feature of Token Ring MAU ports makes Token Ring highly fault-tolerant. For example, if you discover a problematic NIC on the network, you can remove that workstation's cable from the MAU, and the MAU's port will close the ring internally. Similarly, if you discover a faulty MAU, you can remove it from the ring by disconnecting its Ring In and Ring Out cables from its adjacent MAUs and connect the two good MAUs to each other to close the loop.

Now you have learned two definitions for "MAU." Depending on the context, a MAU can refer to a media access unit used as a transceiver in a Thicknet Ethernet network or it can mean a multistation access unit, a hub in a Token Ring network. These are two, unrelated devices. If the term MAU is used on your Network+ certification exam, be certain you understand the context in which it is used before answering the question.

A node on a Token Ring network may also connect to a **Controlled Access Unit (CAU)**. A CAU is a connectivity device similar to a MAU, but in addition to passing data between nodes, a CAU provides more flexibility and easier management of connected nodes. For example, as a network administrator, you could connect to a CAU from your desktop PC and determine what type of traffic is passing through the device or even reconfigure the device. CAUs are more flexible than MAUs because they contain interchangeable modules that you can plug into the Ring In and Ring Out connections. With interchangeable modules, you can use STP for a backbone cable for some time, then upgrade to a fiber-optic backbone by simply inserting the fiber-optic module.

Because of their added functionality, CAUs are more expensive than MAUs. A CAU contains a limited number of receptacles for connected devices. In order to expand the number of nodes you can connect to a CAU, you can plug in a **Lobe Attachment Module (LAM)**. LAMs typically allow up to 20 devices to plug into each CAU receptacle. So, for example, using four LAMs on a single CAU allows 80 devices to be connected to the CAU.

A Token Ring network may use one of three types of connectors on its cables: RJ-45, DB-9, or type 1 IBM. Modern Token Ring networks with UTP cabling use RJ-45 connectors, which are identical to the RJ-45 connector used on 10BaseT or 100BaseT Ethernet networks. Token Ring networks with STP cabling may use a **type 1 IBM connector**, which is depicted in Figure 5-27. Type 1 IBM connectors contain interlocking tabs that snap into an identical connector when one of the connectors is flipped upside-down, making for a secure connection. A **DB-9 connector** (containing 9 pins) is another type of connector found on STP Token Ring networks. This connector is also pictured in Figure 5-27.

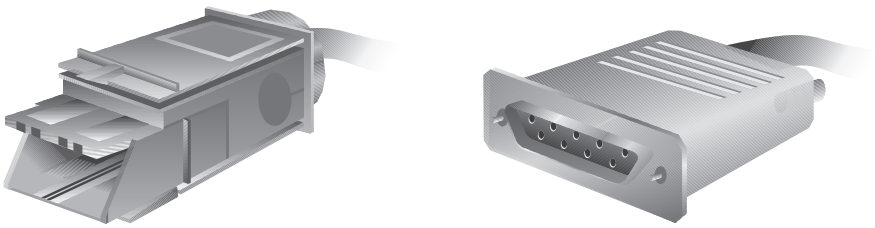

Type 1 IBM connector DB-9 connector

Figure 5-27 Type 1 IBM and DB-9 Token Ring connectors

Occasionally you may work on a network that incorporates a mix of different connectors. (This is more likely to happen on a Token Ring network, given its many connector types.) In this case, you might consider using a **media filter** to enable these different connectors and receptors to fit together. For example, in order to allow a Token Ring NIC with a DB-9 receptor to connect with a network cable that uses an RJ-45 plug, a **Token Ring media filter** is necessary as pictured in Figure 5-28.

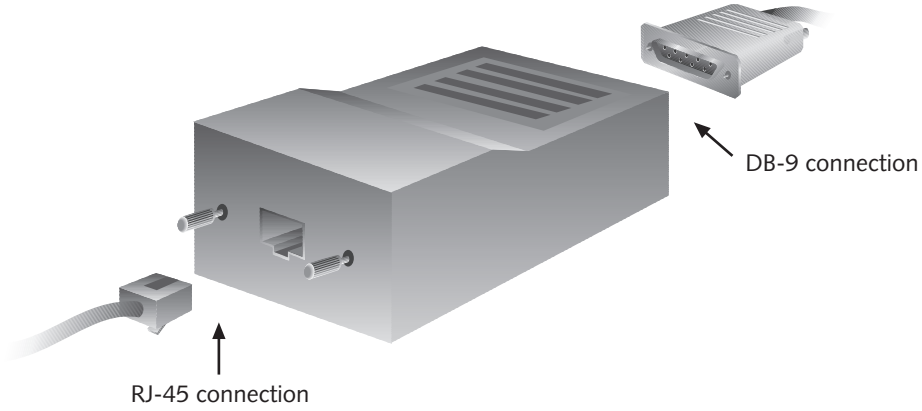

Figure 5-28 A Token Ring media filter

Token Ring Switching

Like Ethernet networks, Token Ring networks can take advantage of switching to better utilize limited bandwidth. Token Ring switching products are typically more expensive and more difficult to manage than Ethernet switches, although they perform essentially the same function. A Token Ring switch can subdivide a large network ring into several smaller network rings. For example, if a 16-Mbps Token Ring network supports 40 users, each workstation has access to approximately 0.4 Mbps. Installing a Token Ring switch that is configured to subdivide the network into four logical subnetworks provides each workstation with approximately 1.6 Mbps (under optimal physical conditions). Thus switching effectively quadruples the bandwidth in this example.

Remember, however, that Token Ring technology does not allow collisions. For this reason, the bandwidth available to each user does not quickly degrade as more users are added; contrast this characteristic to the performance hit that Ethernet takes when more users connect to a single segment.

Token Ring Frames

Token Ring networks may use one of two types of frames: the IEEE 802.5 or the IBM Token Ring frame. The only difference between the two types is that the IBM Token Ring frame adds 2 to 16 octets of routing information that only IBM applications use. Figure 5-29 shows an IEEE 802.5 Token Ring frame.

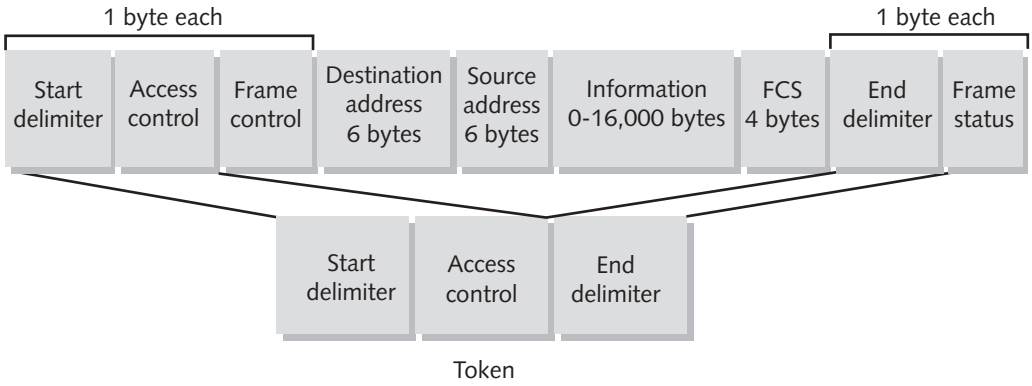

Figure 5-29 An IBM Token Ring frame

Every Token Ring frame includes Starting Delimiter (SD), Access Control (AC), and Ending Delimiter (ED) fields. These three fields of 1 octet each make up the token. Remember that the token is not a frame, but rather is transformed into a frame after a workstation picks it up. The Access Control byte of the token equals 0 bytes if the token is available and 1 if the token is part of a frame currently carrying data, thus signaling that it is not free.

Token Ring frames, just like Ethernet frames, contain destination address and source address fields. As with Ethernet, Token Ring addresses use the MAC address of the device. The destination address is the MAC address of the workstation that will receive the data. The source address is the MAC address of the workstation that transmitted the data. On Token Ring networks, you may also encounter manually administered addresses. Manually assigning network addresses is not a good policy, however. By doing so, you create more work for network administrators and increase the potential for errors.

Both IEEE 802.5 and IBM Token Ring frames contain an information field. Depending on the speed of the Token Ring network, this field can contain from 0 to more than 4000 bytes (for a 4-Mbps network) or from 0 to more than 16,000 bytes (for a 16-Mbps network). Altogether, the maximum frame size for a 4 Mbps network is 4094 bytes; for a 16-Mbps network, it is 17,800 bytes. Notice how much larger the Token Ring frames are than Ethernet frames. As you have learned, larger frame sizes result in more efficient data transmission.

After the data field, each Token Ring frame includes a frame check sequence (FCS). As in Ethernet networks, Token Ring frames use a CRC algorithm to ensure that the data received matches the data sent. The FCS contains the results of this algorithm. In addition, Token Ring frames use a Frame Status (FS) to provide low-level acknowledgment that the frame was received whole.

Design Considerations for Token Ring Networks

If you work on Token Ring networks, you will most likely use a long-established LAN rather than a newly implemented one. In this case, your design considerations will apply to expansion and improvement of an existing infrastructure. Bear in mind these characteristics of Token Ring networks:

- *Cabling*—Token Ring networks can run on shielded or unshielded twisted-pair cabling.
- *Connectivity devices*—Token Ring NICs, switches, hubs, routers, and bridges are generally more expensive than comparable Ethernet equipment.
- *Number of stations*—The number of allowable stations on a Token Ring network is limited, depending on its cabling. You may attach 255 addressable stations on a Token Ring network that runs on shielded twisted-pair, or as many as 72 addressable stations on one that runs on unshielded twisted-pair.
- *Speed*—Token ring networks can run at either 4, 16, or 100 Mbps.
- *Scalability*—You can easily daisy-chain Token Ring MAUs to expand the network. The star layout makes it easy to add nodes to a Token Ring network.
- *Topology*—Token Ring networks are based on a star-wired ring topology, which is highly fault-tolerant.

Fiber Distributed Data Interface (FDDI)

Fiber Distributed Data Interface (FDDI) is a logical topology whose standard was originally specified by ANSI in the mid-1980s and later refined by ISO. FDDI (pronounced "fiddy") uses a double ring of multimode or single mode fiber to transmit data at speeds of 100 Mbps. FDDI was developed in response to the throughput limitations of Ethernet and Token Ring technologies used at the time. In fact, FDDI was the first network transport system to reach the 100-Mbps threshold. For this reason, you will frequently find it supporting network backbones that were installed in the late 1980s and early 1990s. A popular implementation of FDDI involves connecting LANs located in multiple buildings, such as those on college campuses. FDDI links can span distances as large as 62 miles. Because Ethernet and Token Ring technologies have developed faster transmission speeds, FDDI is no longer the much-coveted technology that it was in the 1980s.

Nevertheless, FDDI is a stable technology that offers numerous benefits. Its reliance on fiber-optic cable ensures that FDDI is more reliable and more secure than transmission methods that depend on copper wiring. Another advantage of FDDI is that it works well with Ethernet 100BaseTX technology.

One drawback to FDDI technology is its high cost relative to Fast Ethernet (costing up to 10 times more per switch port than Fast Ethernet). If an organization has FDDI installed, however, it can use the same cabling to upgrade to Fast Ethernet or Gigabit

Ethernet, with only minor differences to consider, such as Ethernet's lower maximum segment length.

FDDI is based on a ring physical topology similar to a Token Ring network, as shown in Figure 5-30. It also relies on the same token-passing routine that Token Ring networks use. However, unlike Token Ring technology, FDDI runs on two complete rings. During normal operation, the primary FDDI ring carries data, while the secondary ring is idle. The secondary ring will assume data transmission responsibilities should the primary ring experience Physical layer problems. This redundancy makes FDDI networks extremely reliable.

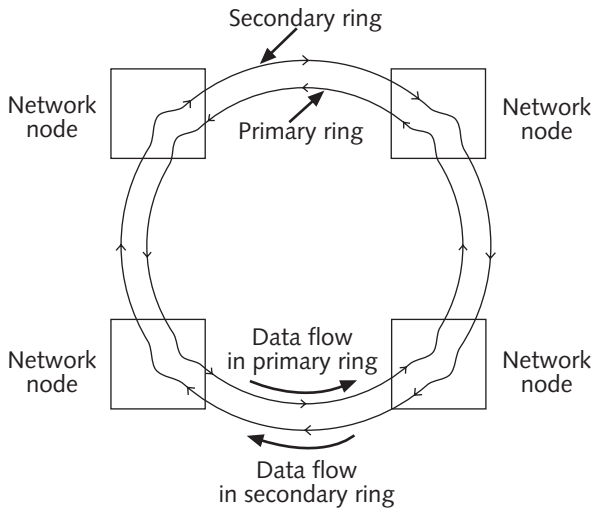

Figure 5-30 A FDDI network

Asynchronous Transfer Mode (ATM)

Asynchronous Transfer Mode (ATM) is a logical topology that relies on a fixed packet size to achieve data transfer rates up to 9953 Mbps. It was first conceived by researchers at Bell Labs in 1983 as a higher-bandwidth alternative to FDDI, but it took a dozen years before standards organizations could reach an agreement on its specifications. ATM may run over specific types of fiber or copper networks, such as SONET or T-carriers (which you will learn about in Chapter 7). It is typically used on WANs, particularly by large data carriers such as telephone companies and Internet service providers.

Like Token Ring and Ethernet, ATM specifies Data Link layer data packaging and Physical layer signaling techniques. But what sets ATM apart from Token Ring and Ethernet is its fixed packet size. The fixed packet in ATM, which is called a **cell**, consists of 48 bytes of data plus a 5-byte header. This fixed packet size allows ATM to provide predictable traffic patterns and better control over bandwidth utilization. However, recall

that a smaller packet size requires more overhead. In fact, ATM's smaller packet size does decrease its potential throughput, but the efficiency of using cells compensates for that loss. Compare ATM's maximum throughput of 9953 Mbps with Fast Ethernet's maximum throughput of 100 Mbps. Even though an ATM cell is a fraction of the size of an Ethernet frame, ATM's throughput is significantly faster.

Another unique aspect of ATM technology is that it relies on virtual circuits. **Virtual circuits** are connections between network nodes that, while based on potentially disparate physical links, logically appear to be direct, dedicated links between those nodes. On an ATM network, switches determine the optimal path between the sender and receiver, then establish this path before the network transmits data. This is an example of circuit switching. In contrast, Ethernet uses packet switching, in which the sending node transmits data first and lets the routers and switches down the wire decide how to direct the data.

The significant benefit to using circuit switching is that it allows ATM to guarantee a specific **quality of service (QoS)**. QoS is a standard that specifies that data will be delivered within a certain period of time after their transmission. ATM networks can supply four QoS levels, from a "best effort" attempt for noncritical data to a guaranteed, real-time transmission for time-sensitive data. This is important for organizations using networks for time-sensitive applications such as video and audio transmissions. For example, a company that wants to use its physical connection between two offices located at opposite sides of a state to carry its voice phone calls might choose the ATM logical topology with the highest possible QoS to carry that data. Without QoS guarantees, data may arrive in the wrong order or too slowly to be properly interpreted by the receiving node.

Because ATM is a recent technology, its developers have made certain it is compatible with other leading network technologies. Its cells can support multiple types of higher-layer protocols, including TCP/IP, AppleTalk, and IPX/SPX. In addition, the ATM logical topology can be integrated with Ethernet or Token Ring networks through the use of **LAN Emulation (LANE)**. LANE encapsulates incoming Ethernet or Token Ring frames, then converts them into ATM cells for transmission over an ATM network.

Currently, ATM is very expensive and, because of its cost, it is rarely used on small LANs and almost never used to connect typical workstations to a network. Gigabit Ethernet—a faster, cheaper, and more standard technology—poses a substantial threat to ATM. In addition to having better-established standards, Gigabit Ethernet is less expensive and a more natural upgrade for the multitude of Fast Ethernet users. It overcomes the QoS issue by simply providing a larger pipe for the greater volume of traffic using the network. While ATM caught on among the very largest carriers in the late 1990s, many networking professionals are now following the Gigabit Ethernet standard rather than spending extra dollars on ATM infrastructure.

Chapter Summary

- A physical topology is the basic physical layout of a network; it does not specify devices, connectivity methods, or addresses on the network. Physical topologies are categorized into three fundamental geometric shapes: bus, ring, and star.
- A bus topology consists of a single cable connecting all nodes on a network without intervening connectivity devices. At either end of a bus network, 50-ohm resistors (terminators) stop signals after they have reached their destination. Without terminators, signals on a bus network experience signal bounce.
- In a ring topology, each node is connected to the two nearest nodes so that the entire network forms a circle. Data are transmitted in one direction around the ring. Each workstation accepts and responds to packets addressed to it, then forwards the other packets to the next workstation in the ring.
- In a star topology, every node on the network is connected through a central device, such as a hub. Any single physical wire on a star network connects only two devices, so a cabling problem will affect only two nodes. Nodes transmit data to the hub, which then retransmits the information to the rest of the network segment where the destination node can pick it up.
- Except for home office networks, few LANs use the simple physical topologies in their pure form. More often, LANs employ a hybrid of more than one simple physical topology.
- The star-wired ring topology is a network that uses the physical layout of a star and the token-passing data transmission method. Data are sent around the star in a circular pattern. Modern Token Ring networks, as specified in IEEE 802.5, use this hybrid topology.
- In a star-wired bus topology, groups of workstations are connected in a star formation; each of the stars is connected to a hub, with all the hubs then networked via a single bus. This design allows you to cover longer distances and easily interconnect or isolate different network segments, although it is more expensive than using either the star or bus topology alone. The star-wired bus topology commonly forms the basis for Ethernet and Fast Ethernet networks.
- Hubs that service star-wired bus or star-wired ring topologies can be daisy-chained to form a more complex hybrid topology.
- A hierarchical hybrid topology can designate hubs at different layers to perform different functions.
- The cabling that connects each hub, or different level of the hierarchy, is called the backbone. A backbone is sometimes called "a network of networks." Backbones usually transmit data at faster speeds than does the cabling that connects each workstation, because they handle the largest loads.
- A serial backbone is the simplest kind of backbone. It consists of two or more hubs connected to each other by a single cable.

- A distributed backbone consists of a number of hubs connected to a series of central hubs or routers in a hierarchy.
- The collapsed backbone topology uses a router or switch as the single central connection point for multiple subnetworks.
- A parallel backbone is the most robust enterprise-wide topology. It is a variation of the collapsed backbone arrangement that consists of more than one connection from the central router or switch to each network segment.
- In a mesh network, routers are interconnected with other routers so that at least two pathways connect each node.
- WAN topologies use the LAN and enterprise-wide topologies as building blocks, but add more complexity because of the distance they must cover, the higher number of users they serve, and heavier traffic they often handle.
- A WAN with single interconnection points for each location is arranged in a peer-to-peer topology. This topology often represents the best solution for organizations with only a few sites and access to dedicated circuits.
- The star topology in a WAN mimics the arrangement of a star LAN. A single site acts as the central connection point for several other points. This arrangement provides several routes for data to follow between any two sites and is, therefore, more reliable than the peer-to-peer or ring WANs.
- In WAN ring topology, each site is connected to two other sites so that the entire WAN forms a ring pattern. This architecture is similar to the LAN ring topology, except that a WAN ring topology connects locations rather than local nodes.
- As with an enterprise-wide mesh, a mesh WAN topology consists of many directly interconnected nodes—in this case, locations. Mesh WANs are the most fault-tolerant WAN configuration. Connecting every node on a network is very expensive, however.
- Tiered WAN topologies are similar to the hierarchical hybrid topologies used with LANs. In a tiered topology, WAN sites connected in star or ring formations are interconnected at different levels, with the interconnection points being organized into layers.
- Network logical topologies (also known as transmission methods) encompass a set of rules specifying which data are packaged and transmitted over network media. The most popular LAN logical topology is Ethernet. Others include Token Ring, LocalTalk, FDDI, and ATM.
- Switching is a component of a network's logical topology that manages the filtering and forwarding of packets between nodes on the network. Every network relies on one of three types of switching: circuit switching, message switching, or packet switching.
- Ethernet is a networking technology originally developed by Xerox in the 1970s and improved by Xerox, Digital Equipment Corporation, and Intel. This flexible technology

can run on a variety of network media and offers excellent throughput at a reasonable cost. Ethernet is by far the most popular logical topology for LANs today.

- Ethernet follows a network access method called Carrier Sense Multiple Access with Collision Detection (CSMA/CD). All Ethernet networks, independent of their speed or frame type, use CSMA/CD.
- On heavily trafficked Ethernet networks, collisions are common. The more nodes that are transmitting data on a network, the more collisions that will take place. When an Ethernet network grows to a particular number of nodes, performance may suffer as a result of collisions.
- A switch is a device that can separate a network segment into smaller segments, each independent of the other and supporting its own traffic. The use of switched Ethernet increases the effective bandwidth of a network segment because fewer workstations vie for the same time on the wire.
- Networks may use one (or a combination) of four kinds of Ethernet data frames: Ethernet IEEE 802.3, Novell Proprietary 802.3, Ethernet II, and IEEE 802.3 SNAP. Each frame type differs slightly in the way it codes and decodes packets of data from one device to another.
- Token Ring networks currently run at either 4 or 16 Mbps, as specified by IEEE 802.5. Token Ring is generally more expensive to implement than Ethernet, but offers high reliability and fault tolerance.
- Token Ring networks use the token-passing routine and a star-ring hybrid physical topology. Workstations connect to the network through Multistation Access Units (MAUs). Token Ring networks may use shielded or unshielded twisted-pair cabling.
- Fiber Distributed Data Interface (FDDI) is a networking standard originally specified by ANSI in the mid-1980s and later refined by ISO. It uses a dual fiber-optic ring to transmit data at speeds of 100 Mbps.
- FDDI's required fiber-optic cable and dual fiber rings offer greater reliability and security than twisted-pair copper wire. It is much more expensive than Fast Ethernet.
- Asynchronous Transfer Mode (ATM) relies on a fixed packet size to achieve data transfer rates up to 9953 Mbps. The fixed packet, called a cell, consists of 48 bytes of data plus a 5-byte header. The fixed packet size allows ATM to provide predictable traffic patterns and better control over bandwidth utilization.
- ATM relies on virtual circuits, logical point-to-point connections that rely on ATM switches, to determine the optimal path between sender and receiver. The ATM switch establishes this path before the network transmits ATM data; in contrast, Ethernet transmits data first and lets the routers and switches down the wire direct the data.
- Applications that benefit from ATM's quality of service guarantees include time-sensitive data, such as video and audio. Currently, ATM is very expensive. It is used on large data carriers' WANs, but rarely on smaller LANs.

KEY TERMS

1 Gigabit Ethernet — An Ethernet standard for networks that achieve 1-Gbps maximum throughput. 1 Gigabit Ethernet runs (preferably) on fiber, but may also run over twisted pair. It is primarily used for network backbones.

10 Gigabit Ethernet — A standard currently being defined by the IEEE 802.3ae committee. 10 Gigabit Ethernet will allow 10-Gbps throughput and will include full-duplexing and multimode fiber requirements.

802.3 Raw — See *Novell proprietary 802.3 frame*.

access method — A network's method of controlling how network nodes access the communications channel. CSMA/CD is the access method used by Ethernet networks.

active monitor — On a Token Ring network, the workstation that maintains timing for token passing, monitors token and frame transmission, detects lost tokens, and corrects problems when a timing error or other disruption occurs. Only one workstation on the ring can act as the active monitor at any given time.

active topology — A topology in which each workstation participates in transmitting data over the network.

Asynchronous Transfer Mode (ATM) — A technology originally conceived in 1983 at Bell Labs, but standardized only in the mid-1990s. It relies on a fixed packet size to achieve data transfer rates up to 9953 Mbps. The fixed packet consists of 48 bytes of data plus a 5-byte header. The fixed packet size allows ATM to provide predictable traffic patterns and better control over bandwidth utilization.

autosense — A feature of modern NICs that enables a NIC to automatically sense what types of frames are running on a network and set itself to that specification.

backbone — The cabling that connects each connectivity device, or the different levels of a hierarchy of connectivity devices.

bus — The single cable connecting all devices in a bus topology.

bus topology — A topology in which a single cable connects all nodes on a network without intervening connectivity devices.

Carrier Sense Multiple Access/Collision Avoidance (CSMA/CA) — A network access method used on LocalTalk networks in which nodes on a shared communication channel signal their intent to transmit data before doing so, thus avoiding collisions.

Carrier Sense Multiple Access/Collision Detection (CSMA/CD) — Rules for communication used by shared Ethernet networks. In CSMA/CD each node waits its turn before transmitting data, to avoid interfering with other nodes' transmissions.

cell — A packet of a fixed size. In ATM technology, a cell consists of 48 bytes of data plus a 5-byte header.

circuit switching — A type of switching in which a connection is established between two network nodes before they begin transmitting data. Bandwidth is dedicated to this connection and remains available until users terminate the communication between the two nodes.

collapsed backbone — A type of enterprise-wide backbone in which a router or switch acts as the single central connection point for multiple subnetworks.

collision — In Ethernet networks, the interference of one network node's data transmission with another network node's data transmission.

collision domain — The portion of an Ethernet network in which collisions will occur if two nodes transmit data at the same time.

Controlled Access Unit (CAU) — A connectivity device used on a Token Ring network. In addition to passing data between nodes, a CAU provides more flexibility and easier management of connected nodes than a MAU.

Cyclical Redundancy Check (CRC) — An algorithm used by the FCS field in Ethernet frames. CRC takes the values of all preceding fields in the frame and generates a unique 4-byte number, the FCS. When the destination node receives the frame, it unscrambles the FCS via CRC and makes sure that the frame's fields match their original organization. If this comparison fails, the receiving node assumes that the frame has been damaged in transit and requests that the source node retransmit the data.

daisy chain — A linked series of devices.

data propagation delay — The length of time data take to travel from one point on the segment to another point. On Ethernet networks, CSMA/CD's collision detection routine cannot operate accurately if the data propagation delay is too long.

DB-9 connector — A connector containing nine pins that is used on STP-based Token Ring networks.

dedicated circuits — Continuous physical or logical connections between two access points that are leased from a communications provider, such as an ISP or local phone company.

demand priority — A method for data transmission used by 100BaseVG Ethernet networks. Each device on a star or hierarchical network sends a request to transmit to the central hub, which grants the requests one at a time. The hub examines incoming data packets, determines the destination node, and forwards the packets to that destination. Because demand priority runs on a star topology, no workstations except the source and destination can "see" the data. Data travel from one device to the hub, then to another device.

distributed backbone — A type of enterprise-wide backbone that consists of a number of hubs connected to a series of central hubs or routers in a hierarchy.

enterprise — An entire organization, including local and remote offices, a mixture of computer systems, and a number of departments. Enterprise-wide computing takes into account the breadth and diversity of a large organization's computer needs.

Ethernet 802.2 frame — See *IEEE 802.3 frame*.

Ethernet 802.3 frame — See *Novell proprietary 802.3 frame*.

Ethernet II frame — The original Ethernet frame type developed by Digital, Intel, and Xerox, before the IEEE began to standardize Ethernet. Ethernet II lacks Logical Link Control layer information but contains a 2-byte type field to identify the upper-layer protocol contained in the frame. It supports TCP/IP, AppleTalk, IPX/SPX, and other higher layer protocols.

Fiber Distributed Data Interface (FDDI) — A networking standard originally specified by ANSI in the mid-1980s and later refined by ISO. FDDI uses a dual fiber-optic ring to transmit data at speeds of 100 Mbps. It was commonly used as a backbone technology in the 1980s and early 1990s, but lost favor as fast Ethernet technologies emerged in the mid-1990s. FDDI provides excellent reliability and security.

Frame Check Sequence (FCS) — A field located at the end of an Ethernet frame that ensures data are received just as they were sent.

hierarchical hybrid topology — A network topology in which devices are divided into separate layers according to their priority or function.

High-Speed Token Ring (HSTR) — A standard for Token Ring networks that operate at 100 Mbps.

hybrid topology — A complex combination of the simple physical topologies.

IEEE 802.3 frame — A popular Ethernet frame type used on IPX/SPX networks. The defining characteristics of its data portion are the source and destination service access points that belong to the Logical Link Control layer, a sublayer of the Data Link layer. Also called LLC or, in Novell lingo, Ethernet 802.2.

IEEE 802.3 SNAP frame — A rarely used Ethernet frame type that is an adaptation of IEEE 802.3 and Ethernet II. SNAP stands for Sub-Network Access Protocol. The SNAP portion of the frame contains the three Logical Link Control fields (DSAP, SSAP, and Control). The Organization ID (OUI) field provides a method of identifying the type of network on which the frame is running. In addition, Ethernet SNAP frames carry Ethernet type information, just as an Ethernet II frame does.

intelligent hub — A hub that, rather than simply regenerating signals, can manage transmissions by dictating which nodes can send and receive data at every instant.

jamming — A part of CSMA/CD in which, upon detecting a collision, a station issues a special 32-bit sequence to indicate to all nodes on an Ethernet segment that its previously transmitted frame has suffered a collision and should be considered faulty.

LAN Emulation (LANE) — A method for transporting Token Ring or Ethernet frames over ATM networks. LANE encapsulates incoming Ethernet or Token Ring frames, then converts them into ATM cells for transmission over an ATM network.

LAN topology — The physical layout, or pattern, of nodes on a local area network (LAN).

layer — A logical division between devices on a network.

LLC frame — See *IEEE 802.3 frame*.

Lobe Attachment Module (LAM) — A device that attaches to a CAU to expand the capacity of that device. LAMs typically allow up to 20 devices to plug into each CAU receptacle.

LocalTalk — A logical topology designed by Apple Computer, Inc. especially for networking Macintosh computers. LocalTalk uses the CSMA/CA network access method, and its throughput is limited to a maximum of 230 Kbps. Because of its throughput limitations, LocalTalk has been replaced by Ethernet on most modern Macintosh-based networks.

logical topology — A networking technology defined by its Data Link layer data packaging and Physical layer signaling techniques. Also known as network transport system or access method.

MacTCP – A version of the TCP/IP protocol supplied with LocalTalk.

media filter — A device that enables two types of cables or connectors to be linked.

mesh network — An enterprise-wide topology in which routers are interconnected with other routers so that at least two pathways connect each node.

mesh WAN topology — A WAN topology that consists of many directly interconnected locations forming a complex mesh.

message switching — A type of switching in which a connection is established between two devices in the connection path; one device transfers data to the second device, then breaks the connection. The information is stored and forwarded from the second device once a connection between that device and a third device on the path is established.

Multistation Access Unit (MAU) — A device on a Token Ring network that regenerates signals; equivalent to a hub.

network access method — See *access method*.

network transport system — See *logical topology*.

Novell proprietary 802.3 frame — The original NetWare Ethernet frame type and the default frame type for networks running NetWare versions lower than 3.12. It supports only the IPX/SPX protocol. Sometimes called 802.3 "raw," because its data portion contains no control bits.

packet switching — A type of switching in which data are broken into packets before they are transported. In packet switching, packets can travel any path on the network to their destination, because each packet contains a destination address and sequencing information.

padding — Bytes added to the data (or information) portion of an Ethernet frame to make sure this field is at least 46 bytes in size. Padding has no effect on the data carried by the frame.

parallel backbone — The most robust enterprise-wide topology. This variation on the collapsed backbone arrangement consists of more than one connection from the central router or switch to each network segment.

peer-to-peer topology — A WAN with single interconnection points for each location.

physical topology — The physical layout of a network. A physical topology depicts a network in broad scope; it does not specify devices, connectivity methods, or addresses on the network. Physical topologies are categorized into three fundamental geometric shapes: bus, ring, and star. These shapes can be mixed to create hybrid topologies.

quality of service (QoS) — The result of standards for delivering data within a certain period of time after their transmission. For example, ATM networks can supply four QoS levels, from a "best effort" attempt for noncritical data to a guaranteed, real-time transmission for time-sensitive data.

ring topology — A network layout in which each node is connected to the two nearest nodes so that the entire network forms a circle. Data are transmitted unidi-

rectionally around the ring. Each workstation accepts and responds to packets addressed to it, then forwards the other packets to the next workstation in the ring.

ring WAN topology — A WAN topology in which each site is connected to two other sites so that the entire WAN forms a ring pattern. This architecture is similar to the LAN ring topology, except that a WAN ring topology connects locations rather than local nodes.

serial backbone — The simplest kind of backbone, consisting of two or more hubs connected to each other by a single cable.

Service Access Point (SAP) — A feature of Ethernet networks that identifies a node or internal process that uses the LLC protocol. Each process between a source and destination node on the network may have a unique SAP.

shared Ethernet — A version of Ethernet in which all the nodes share a common channel and a fixed amount of bandwidth.

star topology — A physical topology in which every node on the network is connected through a central device, such as a hub. Any single physical wire on a star network connects only two devices, so a cabling problem will affect only two nodes. Nodes transmit data to the hub, which then retransmits the data to the rest of the network segment where the destination node can pick it up.

star WAN topology — A WAN topology that mimics the arrangement of star LANs. A single site acts as the central connection point for several other locations.

star-wired bus topology — A hybrid topology in which groups of workstations are connected in a star fashion to hubs that are networked via a single bus.

star-wired ring topology — A hybrid topology that uses the physical layout of a star and the token-passing data transmission method.

switch — The hardware that manages network switching; used to separate a network segment into smaller segments, with each segment being independent of the others, and supporting its own traffic.

switched Ethernet — An Ethernet model that enables multiple nodes to simultaneously transmit and receive data and individually take advantage of more bandwidth because they are assigned separate logical network segments through switching.

switching — A component of a network's logical topology that manages how packets are filtered and forwarded between nodes on the network.

teleconnector — A transceiver used on LocalTalk networks. The teleconnector is linked to the node's serial port on one side, and to the wall jack on the other side.

tiered WAN topology — A WAN topology in which sites are connected in star or ring formations and interconnected at different levels with the interconnection points organized into layers.

token passing — A means of data transmission in which a 3-byte packet, called a token, is passed around the network in a round-robin fashion.

Token Ring media filter — A device that enables a DB-9 cable and a type 1 IBM cable to be connected.

type 1 IBM connector — A type of Token Ring connector that uses interlocking tabs that snap into an identical connector when one is flipped upside-down, making for a secure connection. Type 1 IBM connectors are used on STP-based Token Ring networks.

virtual circuits — Connections between network nodes that, while based on potentially disparate physical links, logically appear to be direct, dedicated links between those nodes.

WAN topology — The physical layout, or pattern, of locations on a wide area network (WAN).

wide area network (WAN) — A network connecting geographically distinct locations, which may or may not belong to the same organization.

Review Questions

1. Under what circumstance might you use a simple bus topology?
   a. when you your LAN services many users
   b. when your LAN services multiple locations
   c. when your LAN services few users
   d. when you want to ensure the highest level of security
   e. when you use hubs to separate workstation groups
2. What kind of topology is susceptible to signal bounce?
   a. mesh
   b. bus
   c. ring
   d. hierarchical
   e. star
3. What are the primary advantages of using a star topology over a ring or bus topology?
4. Most modern networks with more than a few nodes use a hybrid topology. True or False?
5. Why might you want to use a hierarchical topology?
   a. to differentiate levels of connectivity devices and workstation groups
   b. to enable multiprotocol routing between LAN segments
   c. to account for signal bounce between two LAN segments
   d. to use multiple frame types on an Ethernet network
   e. to ensure greater reliability for critical network connections

6. What logical topology, or network transport model, relies most often on a star-wired bus topology?
   a. Ethernet
   b. FDDI
   c. ATM
   d. Token Ring
   e. LocalTalk
7. How do workstations in a ring topology negotiate their data transmissions?
   a. by using CSMA/CD
   b. by using RARP
   c. by using demand priority
   d. by using tokens
   e. by using CSMA/CA
8. Which of the following is a potential problem with daisy-chaining hubs?
   a. exceeding the maximum network length
   b. exceeding the maximum number of workstations per hub
   c. exceeding the maximum collision rate
   d. exceeding the maximum transmission rate
   e. exceeding the maximum number of workstations per segment
9. What type of network backbone is the most reliable?
   a. distributed
   b. collapsed
   c. parallel
   d. serial
   e. hierarchical
10. The Internet is an example of what kind of WAN topology?
    a. peer-to-peer
    b. bus
    c. mesh
    d. ring
    e. tiered

11. Why is packet switching more efficient than circuit switching?
    a. In packet switching, two communicating nodes establish a channel first, then begin transmitting, thus ensuring a reliable connection and eliminating the need to retransmit.
    b. In packet switching, packets can take the quickest route between nodes and arrive independently of when other packets in their data stream arrive.
    c. In packet switching, data are sent to an intermediate node and reassembled before being transmitted, en masse, to the destination node.
    d. In packet switching, packets are synchronized according to a timing mechanism in the switch.
12. Describe the steps a workstation takes under the rules of CSMA/CD.
13. On a 100BaseT (Fast Ethernet) network, what is the maximum number of hubs that can be connected along the bus of a star-wired bus topology?
    a. 2
    b. 3
    c. 4
    d. 5
    e. 6
14. What is the maximum number of addressable stations on a 10BaseT Ethernet network?
    a. 64
    b. 100
    c. 200
    d. 512
    e. 1024
15. Which two of the following might cause excessive data collisions on an Ethernet network?
    a. The network is attempting to use two incompatible frame types.
    b. The overall network length exceeds IEEE 802.3 standards for that network type.
    c. A router on the network is mistakenly forwarding packets to the wrong segment.
    d. A switch on the network has established multiple virtual circuits for a path between two nodes.
    e. A server on the network contains a faulty NIC.

16. What type of media is best suited to 1 Gigabit Ethernet networks?
    a. fiber-optic
    b. unshielded twisted-pair
    c. thick coaxial
    d. shielded twisted-pair
    e. infrared
17. In order to use demand priority on a network (for example, when running 100BaseVG), what type of hub is necessary?
    a. modular
    b. repeater
    c. stackable
    d. intelligent
    e. managed
18. What fields do all Ethernet frame types have in common?
19. At what layer of the OSI Model does framing occur?
    a. Physical layer
    b. Data Link layer
    c. Network layer
    d. Transport layer
    e. Session layer
20. What is the purpose of padding in an Ethernet frame?
    a. ensuring that the frame and data arrive without error
    b. ensuring that the frame arrives in sequence
    c. indicating the length of the frame
    d. indicating the type of higher-layer protocol supported by the frame
    e. ensuring that the data portion of the frame totals at least 46 bytes
21. What is the purpose of a Frame Check Sequence field in an Ethernet frame?
    a. ensuring that data are received without errors at the destination node
    b. ensuring that the frame's length stays constant through transmission
    c. ensuring that the frame is synchronized with other frames in its data stream
    d. ensuring that the frame arrives at the proper destination address
    e. indicating the frame's source address

22. NIC device drivers come with what feature that reduces the need for you to worry about frame types?
    a. autodetect
    b. autonegotiate
    c. autosense
    d. autorespond
    e. autotranslate
23. What are the minimum and maximum sizes for an Ethernet frame?
    a. 46 and 64 bytes
    b. 46 and 128 bytes
    c. 64 and 1518 bytes
    d. 64 and 1600 bytes
    e. 128 and 1600 bytes
24. What type of TCP/IP protocol does LocalTalk use?
    a. MS TCP/IP
    b. MacTCP
    c. AppleTalk
    d. EtherTCP
    e. Apple TCP/IP
25. What is the name of a hub used on a Token Ring network?
    a. Multistation Access Unit
    b. Multiple Carrier Control Unit
    c. Multinode Access Station
    d. Media Access Control Unit
    e. Media Access Unit
26. Which two of the following are disadvantages to using Token Ring networks rather than Ethernet networks?
    a. Their standards are not as well defined as Ethernet's.
    b. They are slower than Ethernet.
    c. They require more expensive connectivity equipment than Ethernet
    d. They are less reliable than Ethernet.
    e. They can't extend as far as Ethernet.

27. Modern Token Ring networks may transmit data at either 4, 16, 32, or 64 Mbps. True or False?
28. If you were working on a Token Ring network that used cables with DB-9 connectors and needed to connect a NIC that contained an RJ-45 receptor, which of the following would help you accomplish your goal?
    a. media access unit
    b. crossover cable
    c. vampire tap
    d. media filter
    e. type 1 IBM connector
29. Which of the following IEEE standards describes Token Ring networks?
    a. IEEE 802.2
    b. IEEE 802.3
    c. IEEE 802.4
    d. IEEE 802.5
    e. IEEE 802.11
30. Which of the following logical topologies is capable of the fastest throughput?
    a. LocalTalk
    b. Fast Ethernet
    c. FDDI
    d. Token Ring
    e. ATM
31. What type of Physical layer is required for FDDI?
    a. a single ring of single mode fiber
    b. a dual ring of single mode or multimode fiber
    c. a dual ring of unshielded twisted-pair cabling
    d. a single ring of shielded twisted-pair cabling
    e. a single ring of multimode fiber
32. Besides their ring-based topologies, what else do FDDI and Token Ring networks have in common?
    a. Both require fiber at the Physical layer.
    b. Both require twisted-pair cabling at the Physical layer.
    c. Both use token passing to mediate data transmission.
    d. Both are less expensive and easier to implement than Ethernet.
    e. Both rely on the parallel backbone structure.

33. You have been asked to serve on a technical committee planning an upgrade from your university's FDDI network to a Gigabit Ethernet network. The rest of the committee asserts this will be a relatively simple transition. What concern should you raise that contradicts their assertion?
    a. All of the FDDI fiber will have to be dug up and replaced with single mode fiber.
    b. The maximum allowable distance for a FDDI network is longer than that of an Ethernet network, so the existing FDDI network will need to be divided into smaller subnetworks.
    c. The connectors on the ends of the FDDI cables are SMA connectors, which will not fit into the receptors on Gigabit Ethernet routers.
    d. Since the transmission rate on the Ethernet network will be less than that of the FDDI network, users will notice dramatically slower response times.
34. What type of switching do ATM networks use?
    a. circuit switching
    b. packet switching
    c. multiprotocol switching
    d. message switching
    e. Layer 1 switching
35. Which two of the following might explain why network administrators prefer Gigabit Ethernet over ATM?
    a. Gigabit Ethernet is a more natural upgrade for their existing Ethernet networks.
    b. Gigabit Ethernet is typically less expensive to implement than ATM.
    c. Gigabit Ethernet can carry TCP/IP traffic, while ATM cannot.
    d. Gigabit Ethernet is endorsed by Microsoft and Novell, while ATM is not.
    e. Gigabit Ethernet offers quality of service guarantees, while ATM does not.

HANDS-ON PROJECTS

Project 5-1

In this exercise, you will create a simple star-wired bus network, one of the most typical forms of an Ethernet network. This project requires two Ethernet 10-Mbps hubs, containing at least four ports each, six (straight through) patch cables, three workstations, and one file server, all with 10-Mbps NICs (installed and correctly configured). The server should be running Windows 2000, and the workstations should run either Windows 98 or Windows 2000. All should have TCP/IP properly installed and configured. Make sure

that user accounts are established that can be used to log onto the server. Finally, you will also need a paper and pencil.

1. Make sure that all the hubs, workstations, and the server are plugged in.
2. Connect the two hubs to each other by inserting one end of a patch cable in one hub's link port and the other end of the patch cable in the second hub's link port. Turn on both hubs if they are not already on.
3. Using another patch cable, connect one of the workstations to another port in the first hub. In the same manner, connect the server to the first hub, then turn on the workstation and server. Notice what happens to the lights on the hub when the workstation and server start up.
4. Repeat Step 3, but connect two different workstations to the second hub and then turn on the workstations.
5. Log onto the server from one of the four workstations. If you can see the server's resources, you have successfully created a star-wired bus Ethernet network, where the two hubs form the network's backbone. If you cannot log onto the server, check the cable connections from your workstations to the hub, between the hubs, and between the hub and the server.
6. On a separate piece of paper, draw the physical topology you have just created, marking the server, workstations, hubs, and cables on your drawing.

Project 5-2

In this exercise, you will use Windows 2000 Server's Network Monitor. Network Monitor is a tool that comes with the Windows 2000 Server operating system that allows you to view different data frames traveling to and from a server's NIC. You can also use it to determine various characteristics about the frames, including the networking protocols they carry. Network Monitor is especially useful when troubleshooting network connection problems. This project requires a Windows 2000 server that is running IPX/SPX and TCP/IP, with at least two clients attached and logged in. (If a Windows 2000 server is not available, a Windows NT server with the Network Monitor tool installed will also work.) The clients may be Windows 98 or Windows 2000 workstations.

1. Log onto the Windows 2000 server as an administrator.
2. On the Windows 2000 server, click **Start**, point to **Programs**, point to **Administrative Tools**, then click **Network Monitor**. The Network Monitor window opens. (Maximize the Network Monitor screen if it does not maximize automatically.)
3. If your server is connected to more than one network, you may be asked to select the network on which you want to monitor data. Choose the network to which your clients are connected (the local network). In the next step, you will begin the capture process to gather data that can later be analyzed.

4. Click **Capture** on the menu bar, then click **Start** to begin capturing network traffic information. In the next step, you will generate traffic to and from the server by accessing its shared resources.
5. From one of the workstations that is logged into the server, open a file on the server (for example, a spreadsheet or word-processing document).
6. From another workstation that is logged into the server, double-click the **Network Neighborhood** icon (in the case of a Windows 98 workstation) or **My Network Places** icon (in the case of a Windows 2000 workstation). Find the server's icon and double-click that. Then open some shared folders on the server.
7. Now that you have generated traffic to and from the server, return to the server to stop capturing data. In the Network Monitor menu bar, click **Capture**, then click **Stop**.
8. In the Network Monitor menu bar, click **Capture**, then click **Display Captured Data**.
9. As shown in Figure 5-31, at the bottom of the screen, Network Monitor lists the packets it has captured. Double-click one of the packets to view more information about it.

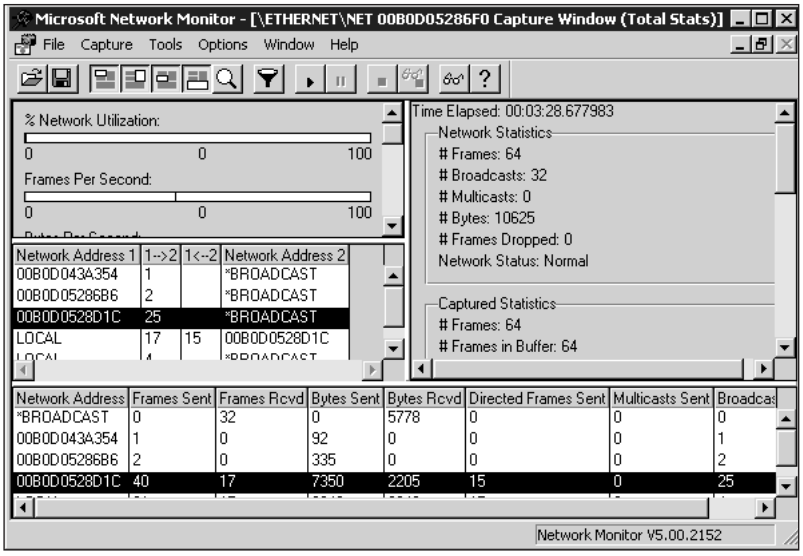

Figure 5-31 The Windows 2000 Network Monitor window

10. The Nework Monitor window displays three panes. The top pane contains a list of captured frames; the middle frame provides header and delivery details; and the bottom pane shows a hexadecimal/ASCII representation of the frame's content.

11. In the top pane, choose to view a frame that has a different kind of protocol listed in its protocol column than the frame you just viewed. Browse through its fields to find its protocol type, length, and source address information.
12. To close Network Monitor, click **File**, then click **Exit**. Do not save the information you have captured.

Project 5-3

In this exercise, you will survey organizations in your area to determine which network transport model, cabling types, and transmission speeds are used on their LANs. From this information, you can determine the most popular networking design approaches. You may also be able to tell what approaches will grow in popularity and which ones will become obsolete.

1. Identify five businesses, schools, or civic organizations in your area that use networking technology—for example, an insurance company, a utility company, a local school district, a chain of retail stores, an architectural firm, or an ISP.
2. For each organization, find contact information for its manager of the Information Technology (IT) department.
3. Call the IT department manager and ask him or her the following questions:
   - Does your network use Ethernet, Token Ring, or both?
   - What transmission speed does your network use?
   - On what type of backbone does your network rely?
   - What type of wiring does your network use? Do you use the same type of wiring for the backbone as you do for connecting workstations?
   - Do you use WAN technologies? If so, in what kind of a topology are your separate locations arranged?
   - If you had all the money you wanted in your budget, what type of upgrades would you make to your network?
   - If you had money enough to make only one of these upgrades, which one would it be?
4. Compile the answers from all five managers. Which network transport method is most popular? Which transmission speed? Which type of backbone? How do the economics of each networking approach affect the manager's decisions regarding wiring and logical topology? Compare your results with the results of others in your group.
5. Write and send a thank you note to each of the IT managers you interviewed, thanking them for sharing their valuable time with you.

CASE PROJECTS

1. You have been asked to design a LAN for a very successful CPA firm with five departments in one building and a total of 560 employees. Currently, the firm has no networked computers, and it is open to any suggestions you can offer. The firm does have a few requirements, however. It wants to make sure that it can easily expand its LAN in the future without exorbitant costs and moving a lot of equipment. The firm also wants to make sure that every department has very fast access to the LAN, and, of course, it wants the LAN to remain up at all times. It has already decided to use the NetWare 5.0 network operating system. What kind of LAN will you design for this company?
2. AstroTech Components, a company that manufactures parts for the aeronautics industry, is having trouble with a network segment in one of its plants and has asked you for help. According to the network administrator, the plant was incorporated into the existing Ethernet 10BaseT network two weeks ago. Since then, the users have been complaining of intermittent lockups, software errors, and disconnections. During your visit, she shows you the very organized telecommunications closet. Then, the network administrator escorts you to the production floor, where she points out the 20 Windows 98 machines that the supervisors use to enter numbers into a database from their desks in the production area. The supervisors try to explain their problems in detail, but you and the network administrator can barely hear above the roar of stamping machines. You begin to walk away when you notice that the network cabling is strung along the outside of support posts between stamping machines. When you reach her office, what suggestions do you make to the network administrator to fix the problem?
3. Because you solved AstroTech's plant dilemma so quickly, the network administrator has time to ask you more questions. In particular, she is concerned about the CAD/CAM workgroup. The users in this group fought for months to get new machines. Now that her technicians finally installed the more powerful workstations, however, the users can't access the network. You ask what kind of network they are on, and the network administrator says that this group was upgraded to 100BaseT, along with two other groups, just yesterday, because these users needed the extra speed. When you ask whether all users are affected, she says that everyone—even the department vice president, who has full rights to the network—is prevented from logging on. You suspect that the CAD/CAM users' network access is the problem. What steps do you take next?

4. The network administrator understands everything you've explained so far, and although your solution will cost a little more, she's glad to have the company's CAD/CAM workgroup problem resolved. Now she asks about the Finance Department, which is experiencing problems logging on. These personnel are running on Windows 95 workstations connected to a 10BaseT bus Ethernet segment. The problem happens about half the time. Once they're logged on, these users occasionally experience other problems, but they haven't recorded any of the error messages. For a long time, the network administrator thought that Finance staff members were just forgetting their passwords, but her technicians have verified that the connectivity problem is real. She admits that the wiring closet for the Finance area is a mess because the department doubled in size when the company bought out its main competitor, Solstice, Inc., a few months ago. What do you think might be causing the problem?
5. Before you leave, the network administrator asks your opinion about upgrading the rest of the company to 100BaseT from 10BaseT. Although her technicians have told her that this move is necessary, she is concerned about the costs associated with replacing wiring, NICs, switches, and hubs. As it is, she has to purchase 500 new desktop PCs this year. The network administrator has heard about 1 Gigabit Ethernet and wonders whether it would be better to wait for that architecture. What considerations do you point out that might help her with her decision?